Quantum States: What the Hell Are They?

The Post-Växjö Phase Transition

Knowledge $\longrightarrow$ *Information* $\longrightarrow$ *Belief* $\longrightarrow$ *Pragmatic Commitment*

Christopher A. Fuchs

Computing Science Research Center
Bell Labs, Lucent Technologies
Room 2C-420, 600–700 Mountain Ave.
Murray Hill, New Jersey 07974, U.S.A.

28 June 2001 — 29 June 2002

Abstract

Collecting it up, it's hard to believe I've written this much in the little time since Växjö. I guess it's been an active time for me. I think there's no doubt that I've gone through a phase transition. For all my Bayesian rhetoric in the last few years, I simply had not realized the immense implications of holding fast to the view that "probabilities are subjective degrees of belief." Of course, one way to look at this revelation is that it is a *reductio ad absurdum* for the whole point of view—and that will certainly be the first thing the critics pick up on. But—you wouldn't have guessed less—I'm starting to view it as a godsend. For with this simple train of logic, one can immediately stamp out the potential reality/objectivity of any of a number of terms that might have clouded our vision. With so much dead weight removed, the little part left behind may finally have the strength to support an ontology.

— from a letter to Mermin, 9 October 2001

Contents

> Now my own suspicion is that the Universe is not only queerer than we suppose, but queerer than we can suppose.
>
> — J. B. S. Haldane

28 June 2001, to Khrennikov, "Context Dependent Probability"

Khrennikovism 1 : *It was nice to meet you in Växjö and discuss fundamental problems of quantum theory. Unfortunetely, I have the impression that my presentation on Contextual Probabilistic Interpretation of quantum theory was not so clear for participants (conversations during lunches and dinners). I try to present my views as short and clear as possible.*

Thank you for valuing my opinion on your ideas; I am flattered. So I treated the problem in a conscientious manner: I downloaded and read three of your papers (`quant-ph/0103065`, `0105059`, and `0106073`).

I am indeed quite intrigued by the possibility that quantum mechanics may be nothing more than a calculus for comparing probabilities when the experimental context cannot be deleted from the results it brings about. In vague philosophical terms, I think this is precisely the kind of idea Bohr, Heisenberg, and Pauli were bandying about in constructing their interpretation of quantum mechanics. It is certainly the kind of notion Bohr was trying to get at with his emphasis on "complementarity." So I would welcome a more precise way (a mathematical way) of expressing the essence of all this. I myself have been attracted to this sort of thing for a long time: it is a large part of the thread connecting my "Notes on a Paulian Idea"—that is, that the observer sets the context, and, in the words of Pauli, cannot be "detached" from what he finds. Also you can find discussions of it in Sections 4 and 8 of the large paper I was circulating at the conference, "Quantum Foundations in the Light of Quantum Information." I say all this to make it clear that I am more than sympathetic to your program.

However, as much as I would like to tell you otherwise (because you are my friend), I do not see that your present formulation of the problem moves very far toward quantum mechanics in a convincing way. There are problems on at least two levels.

Maybe the most devastating and immediate is your move between Eqs. (5) and (6) of `quant-ph/0106073`. (I'll focus on that paper for specificity since I did not see you make a stronger argument in either of the other two papers.) You write:

> The perturbation term $\delta(\mathcal{S}, \mathcal{S}')$ depends on absolute magnitudes of probabilities. It would be natural to introduce normalized coefficient of the context transition ...

The question anyone will ask is, "Why is this natural?" What compels the precise form of the normalization other than that it forces the equation to look of a more quantum mechanical form. Why did you choose the square root rather than the third root, say? Indeed, why not divide by the absolute value of δ, or the exponential of δ, or any other combination of functions one could pull out of a hat? To put it not so gently, it looks as if you built the desired answer in at the outset, with little justification otherwise.

The second level of my problem is that, even if you do get this far, how do you make the further step to vector space representations of quantum mechanics? Why are observables POVMs and not other exotic entities? What leads us to the starting point of Gleason's theorem? Etc., etc.? I don't

see that you have enough structure to do that. But more importantly, until you have done that I would have to say that your theory remains fairly empty in making a connection to quantum mechanics. Too empty.

The way I view the problem presently is that, indeed, quantum theory is a theory of contextual probabilities. This much we agree on: within each context, quantum probabilities are nothing more than standard Kolmogorovian probabilities. But the contexts are set by the structure of the Positive Operator-Valued Measures: one experimental context, one POVM. The glue that pastes the POVMs together into a unified Hilbert space is Gleason's "noncontextuality assumption": where two POVMs overlap, the probability assignments for those outcomes must not depend upon the context. Putting those two ideas together, one derives the structure of the quantum state. The quantum state (uniquely) specifies a *compendium* of probabilities, one for each context. And thus there are transformation rules for deriving probabilities in one context from another. This has the flavor of your program. But getting to that starting point from more general considerations—as you would like to do (I think)—is the challenge I haven't yet seen fulfilled.

I very much hope that I have not offended you with these comments. I greatly respect your program. But because of that I want much from it. I want it to stretch our understanding. John Wheeler used to say, "We must make as many mistakes as we can, as fast as we can, or we'll never have a hope of gaining a true understanding!" I let that philosophy rule my research life. Thus I can only commend you for your exploration, and hold the strongest hope that something firm will come from it with a little more work and contemplation.

28 June 2001, to Mermin, "The Oblique Observer"

You haunt me so that I wake up in the middle of the night just to spar with you. It's this damned noncontextuality of quantum probabilities. If one walks into the game with the *firm* belief that a quantum state is a state of knowledge, then noncontextuality is almost a given. That, in part, is what my section "Whither Bayes' Rule?" is about. Moreover, all questions about instantaneous signaling through quantum-state change just become silly: such questions spring solely from a wrong-headed view.

But, you do not walk into this game with the firm belief that a quantum state is a state of knowledge. So I am forced to work to win your approval. In the end, with clarity achieved, your demands will have been a great gift. Right now, they are just annoying.

In any case, this leads me down the path of the oblique observer. There is a way in which a von Neumann-like view of measurement may be a virtue. That is, a view of measurement where, to make it go, I first introduce an ancilla to interact with the system. Then I introduce a second ancilla to interact with the first ancilla because I didn't know how to solve the measurement problem at the new level. Then I introduce a third ancilla, and so forth. Von Neumann went to this extreme because he was chicken to let the mental update that is a measurement fall outside of physics (as Thomas Bayes would have). So, he piled up superobserver after superobserver, just so there would always be an outside view. But so be it: Let us glean what we can from this freedom of our descriptions.

The main point is this: You pick any measurement (any POVM) you wish for some system, and I can always think of a way to get at that measurement in an indirect way. That is to say, I can always delay my cutting of the Gordian knot until I get to a system with no residual causal link to the one I'm really interested in. In language I don't like, it means I can always induce the "physical collapse" somewhere else. I can always push the measurement to an arena where you would take the noncontextuality assumption (on the system of interest) as a given. My direct measurement on

an ancilla serves only to refine my knowledge about the actual system: the actual system cannot care how I came to that refinement, or, indeed, if I ever pursued refining my knowledge at all.

You get my point—with that much repetition, you'd better. This gives rise to a vague idea that perhaps you can help me elicit into reality. (And save me some torture that you are the root of.) Forgetting about the precise structure of quantum mechanics, why should we not view all observations as oblique observations? Whose philosophy ever dreamt that we had *direct* access to the minutest details of the world in the first place? (I'll paste in the great Heisenberg/Einstein quote below so you'll have quick access in case you wish to remind yourself of it—especially the fourth, fifth, and sixth sentences.)

Moreover, coming back to quantum mechanics in particular, what is to keep me (in my derivation of the tensor product rule) from thinking of the two separate observations as each concerning the opposite system? The question is, can one get some quantitative mileage from this. At least that's the question on my mind as the sun is rising. Any thoughts?

It's a good thing you left the Växjö before Friday's lunch. The boiled potatoes were ridden with rocks: I lost an eighth of a molar in the process! What was your overall impression of the meeting (good and bad)?

28 June 2001, to Mermin, "The Kettle Black"

I'm just reading through a friend's comments on my NATO draft. Just after my equations 65–67, where I write, "The resemblance between the process in Eq. (66) and the classical Bayes' rule of Eq. (38) is unmistakable," he writes: "Seems to me contrived — you want to find a resemblance and then you find it."

Thinking back to one of his recent productions, I think "Boy, that's the pot calling the kettle black!"

But aside from that, I do seem to be impressing less people with this result than I would have thought. Is it my presentation, or is it really the substance? This result, in particular, is the one I'm most pleased with in the paper, but it gets the coldest reception. You've never yet told me your truest thoughts on that section. I'd like to hear them.

02 July 2001, to Chakalov, "Objective Properties"

Thank you for all the interest you've shown in the papers I have been involved with. I commend you in your efforts to get to the bottom of what's going on in our world. But I cannot believe it very likely that distinct new kinds of *physics* arise in our brain processes. Instead the road I have chosen to develop is making sense of quantum mechanics (as a theory predominantly of inference) from *within* quantum mechanics. I understand that your road is distinct: but life is short, and one has to make a cut or one will certainly never get anywhere. My own direction may turn out to be completely wrong, but I have decided to pursue it with dogged determination and not to get derailed.

I wish you luck in your own pursuits.

02 July 2001, to Finkelstein, "Making Good Sense"

Finkelsteinism 1: *Would you agree?*

Thanks for the note! Yes, I guess I would (though only to the small extent that I think many-worlds is coherent in the first place). But there are two things working in the background. 1) Probably plain sloppiness on our part in our wording. And 2) the fact that Ruediger is (presently) more conciliatory to MWI than Carl and I are. He sees Bayesian probability as holding a place even in their interpretation. (A rough cartoon is: In their interpretation, the universal wave function serves an ontologic role, while the relative states in a Schmidt decomposition with respect to an observer's mind serve the same epistemic role we ascribe to them.) We should probably either remedy 1) or make 2) more clear, or both. We'll have to huddle for that.

Finkelsteinism 2: *(And also by pointing out that the experiment with detectors in relative motion was reported in Zbinden et al (*`quant-ph/0007009`*), rather than in your ref. 2.)*

That's probably my screw-up: I just assumed (from a search through SciSearch) that Ref. 2 was the published version of Zbinden et al. I presume you're telling me it's not. Is there some history here that we should be aware of? Or a different published reference?

Finkelsteinism 3: *I enjoyed reading your latest "quantum states are states of knowledge" manifesto,*

And there's still another one coming: it should have appeared on `quant-ph` today. I hope you'll read it too. This one's a solo flight by me (titled "Quantum Foundations in the Light of Quantum Information"). (BTW: Don't let the sober sounding abstract on `quant-ph` fool you; I'm as loquacious and philosophical as usual on the inside.)

02 July 2001, to Schack, "quant-ph/0106133"

I would only temper what you just said by making one addition:

Schackcosm 1: *Actually, I believe that the relative state an observer in some branch of a multiverse* [has no choice but to assign] *to, say, a qubit has a very natural interpretation as a state of knowledge.*

The lack of free choice is important there, and to that extent the whole scheme is non-Bayesian. Bayesian probabilities are never fixed by edict. In a way, this is just a fancy version of David Lewis's principal principle.

Top of the mornin',

04 July 2001, to Khrennikov, "Re: Invitation"

Khrennikovism 2: *Yes, this is very well! However, for me, the only bridge between "reality" and our subjective description is given by relative frequencies ...*

But there other ways to make the bridge: this is what gambling situations (like the Dutch-book argument that Schack spoke about) are about. They give a NON-frequency OPERATIONAL definition to probabilities. Subjective probabilities make their OBJECTIVE mark on the world by specifying how an agent should act when confronted with them.

04 July 2001, to Khrennikov, "Re: Context Dependent Probability"

Khrennikovism 3: *P.S. But! How can you unify contextuality with subjective probability?*

I just don't see this as a problem. In choosing one experiment over another, I choose one context over another. The experiment elicits the world to do something. To say that the world is indeterministic means simply that I cannot predict with certainty what it will do in response to my action. Instead, I say what I can in the form of a probability assignment. My probability assignment comes about from the information available to me (how the system reacted in other contexts, etc., etc.). Similarly for you, even though your information may not be the same as mine. The OBJECTIVE content of the probability assignment comes from the fact that NO ONE can make *tighter* predictions for the outcomes of experiments than specified by the quantum mechanical laws. Or to say it still another way, it is the very existence of transformation RULES from one context to another that expresses an objective content for the theory. Those rules apply to me as well as to you, even though our probability assignments WITHIN each context may be completely different (because they are subjective). But, if one of us follows the proper transformation rules—the quantum rules—for going to one context from another, while the other of us does not, then one of us will be able to take advantage of the other in a gambling match. The one of us that ignores the structure of the world will be bitten by it!

05 July 2001, to Finkelstein, "Standing Up and Saying YES"

Thanks for the comments. I welcome any that you send me!

Finkelsteinism 4: *This is not really any objection to what you have written, but the story you tell on page 10 might produce even WEAKER knees with the following modification: Suppose that Alice, instead of choosing ANY state $|\psi\rangle$ for her qubit, makes her choice from a finite and previously-agreed-upon set. She broadcasts the result of her measurement, but keeps her choice a secret, except that she reveals her choice in a sealed envelope which she sends to Chris (who initially leaves it sealed). Bob performs the appropriate Pauli rotation, then he makes a guess as to which state Alice chose, and performs a yes-no measurement with that guess; he communicates his guess, as well as the yes-no result, to Chris.*

Chris can now open the sealed envelope; if it happens that Bob's guess was in fact correct, then the result must have been "yes". So, if one wanted to be contrary (and of course I do not) one might say that, although when the yes-no measurement was performed nobody knew that the guess was correct, and although Alice did not "take the time to ... interact with it", nevertheless the qubit had "the power to stand up and say YES all by itself".

I agree, this does sound even more dramatic. And maybe I will start using it in my presentations. But the point remains the same: it is Bob's action that elicits a consequence.

You can see, I keep dreaming (modern) alchemical thoughts. Below. From: W. Heisenberg, "Wolfgang Pauli's Philosophical Outlook," in his book *Across the Frontiers*, translated by P. Heath, (Harper & Row, New York, 1974), pp. 30–38.

> The elaboration of Plato's thought had led, in neo-Platonism and Christianity, to a position where matter was characterized as void of Ideas. Hence, since the intelligible was identical with the good, matter was identified as evil. But in the new science the

> world-soul was finally replaced by the abstract mathematical law of nature. Against this one-sidedly spiritualizing tendency the alchemistical philosophy, championed here by Fludd, represents a certain counterpoise. In the alchemistic view "there dwells in matter a spirit awaiting release. The alchemist in his laboratory is constantly involved in nature's course, in such wise that the real or supposed chemical reactions in the retort are mystically identified with the psychic processes in himself, and are called by the same names. The release of the substance by the man who transmutes it, which culminates in the production of the philosopher's stone, is seen by the alchemist, in light of the mystical correspondence of macrocosmos and microcosmos, as identical with the saving transformation of the man by the work, which succeeds only 'Deo concedente.'" The governing symbol for this magical view of nature is the quaternary number, the so-called "tetractys" of the Pythagoreans, which is put together out of two polarities. The division is correlated with the dark side of the world (matter, the Devil), and the magical view of nature also embraces this dark region.

and

> When, in the spring of 1927, opinions on the interpretation of quantum mechanics were taking on rational shape and Bohr was forging the concept of complementarity, Pauli was one of the first physicists to decide unreservedly for the new possibility of interpretation. The characteristic feature of this interpretation—namely, that in every experiment, every incursion into nature, we have the choice of which aspect of nature we want to make visible, but that we simultaneously must sacrifice, in that we must forego other such aspects—this coupling of "choice and sacrifice," proved spontaneously congenial to Pauli's philosophical outlook. In the center of his philosophical thinking here there was always the wish for a unitary understanding of the world, a unity incorporating the tension of opposites, and he hailed the interpretation of quantum theory as a new way of thinking, in which the unity can perhaps be more easily expressed than before. In the alchemistic philosophy, he had been captivated by the attempt to speak of material and psychical processes in the same language. Pauli came to think that in the abstract territory traversed by modern atomic physics and modern psychology such a language could once more be attempted ...

05 July 2001, to Khrennikov, "Re: Invitation"

Khrennikovism 4: *I think you (and everybody) do in the following way: you have some experience with gambling (frequency!) and use this experience to introduce "subjective" probabilities.*

No, I think it is just the opposite: people almost never use frequency data as the determiners of their information in any common situation. Instead they use symmetry. If someone presents me with a coin that I have never seen before, then after a quick examination, I will likely ASCRIBE a 50/50 probability to its coming up heads SIMPLY because I have no reason to believe otherwise. But if Danny Greenberger is the tosser of it, I know that he has the skill to make it look superficially as if it were being tossed in a haphazard fashion but it will still come up heads every time. The 50/50 ascription is not a property of the coin! It is simply a property of ignorance.

09 July 2001, to Mermin, "The O'bleak Observer"

Merminition 1: *You should not dismiss my feeling that you've not adequately justified your assumption about noncontextuality as merely a manifestation of a regretable atavistic tendency to reify the quantum state.*

The "oblique observer" note was a concession, not a dismission. It is evidence that I am taking your point very seriously (even though I'd rather being out playing with the other kids).

Merminition 2: *The question you're evading is what it means for one and the same positive operator E to appear in many different POVMS.*

No, I don't think I'm evading it. It means that those various interventions or ways of gathering data—those POVMs—physically diverse though they may be—all lead to at least one common possibility for what my knowledge can be updated to (modulo the unitary readjustment).

09 July 2001, to Mermin, "More O'bleakness"

Merminition 3:

> CAF Said: Moreover, all questions about instantaneous signaling through quantum-state change just become silly: such questions spring solely from a wrong-headed view.

No! The signalling has nothing to do with quantum-state change. (We've been through this before.) If Bob and Alice share a large number of identically prepared pairs, then a very reasonable requirement is that the statistical distribution of outcomes Bob gets from his members of the pairs cannot depend on what Alice chooses to do to her members. (If it did Alice could send useable unmediated signals to Bob.) Again, this has nothing to do with how you like to think about probabilities or quantum states. I offered this to you as an example of a situation in which you can, in fact, justify the non-contextuality of certain probabilities by appealing to an independent physical requirement (no remote signalling).

I hold firm in my opinion. It has *everything* to do with how you like to think about (quantum) probabilities. If you think the probabilities are subjective expectations for the local consequences of one's experimental interventions, then the question never arises.

However, granting you a little distrust for that, the point about oblique observations is that one might always be able to think of a quantum measurement as being enacted on a system other than the intended one. This would give your point above a natural means for being used to justify noncontextuality for *all* quantum measurements.

Again, I'm starting to feel awfully comfortable with noncontextuality as the very simplest generalization of Bayes' noncontextuality. It is the very glue that puts measurement outcomes into Hilbert space in the first place. (Otherwise we might just draw out an exhaustive list of one-outcomed, two-outcomed, three-outcomed measurements etc., etc., and never even suppose a connection between them.) But I offer the above as an effort to go in the direction you want me to.

10 July 2001, to Landahl, "Replies on a Preskillian Meeting"

Wow, what a set of notes! Thank you all for the interest in my silly efforts. This is a little unexpected.

Landahlism 1: *The talk was a smash. It went much better than I was expecting, causing much discussion. (As you know, some of the people in our group are rather reticent, so that's really saying something.)*

I am so glad to hear that. It is really very flattering.

Landahlism 2: *At the end of the talk, Sumit decided to go up to the chalkboard and take a poll, the topic and results of which I'll leave as a surprise, as I imagine John will tell you about them himself. (If he doesn't, just e-mail me back and I'll let you know.)*

I presume this is the poll Scott mentioned. I'll say more about that later.

Landahlism 3: *I also mentioned your program to establish an information-theoretic foundation for all the laws of quantum mechanics (and physics?) in this section. This proposal met with much skepticism from the audience. I'm somewhat sympathetic to your cause (certainly more so than some of our denizens!), but I don't believe that all of physics has an information theoretic description. For example, where would the (dimensionful) physical constants enter into this scheme (like Planck's constant and the speed of light)? I don't see how they could enter unless they define what physical dimensions are, which is rather peculiar.*

I really am very flattered by all this attention, but I do get dismayed when I can't seem to get the most important point across to my readers. Even sympathetic readers! Because of this, I have spent months and months trying to clarify and refine my presentation. But for some reason it is amazingly difficult to get the point across. At the very least I need people to understand what I want *before* they declare that they disagree with it. (Disagreeing with it would *then* be fair enough.) The sentences above seem to convey that you haven't gotten to the level of understanding what I want. How can the following sentences be consistent with what you say above?

1. Abstract, penultimate sentence.

   This method holds promise precisely because a large part (but not all) of the structure of quantum theory has always concerned information.

2. Section 1, last paragraph.

   Our foremost task should be to go to each and every axiom of quantum theory and give it an information theoretic justification if we can. Only when we are finished picking off all the terms (or combinations of terms) that can be interpreted as information—subjective information—will we be in a position to make real progress. The raw distillate that is left behind, miniscule though it may be, will be our first glimpse of what quantum mechanics is trying to tell us about nature itself.

3. Section 2, last paragraph.

   The world is sensitive to our touch. ... The whole structure of quantum mechanics—*it is speculated*—may be nothing more than the optimal method of reasoning and processing information in the light of such a fundamental (wonderful) sensitivity.

4. Section 3, penultimate paragraph, page 9.

   The complete disconnectedness of the quantum-state change rule from anything to do with spacetime considerations is telling us something deep: The quantum state is information. Subjective, incomplete information. Put in the right mindset, this is *not* so intolerable. It is a statement about our world. There is something about the world that keeps us from ever getting more information than can be captured through the formal structure of quantum mechanics. Einstein had wanted us to look further—to find out how the incomplete information could be completed—but perhaps the real question is, "Why can it *not* be completed?"

5. Section 5, last two paragraphs.

   Perhaps the structure of the theory denotes the optimal way to reason and make decisions in light of *some* fundamental situation, waiting to be ferreted out in a more satisfactory fashion.

   This much we know: That "fundamental situation"—whatever it is—must be an ingredient Bayesian probability theory does not have. There must be something to drive a wedge between the two theories. Probability theory alone is too general of a structure. Narrowing it will require input from the world about us.

6. Section 7, last two paragraphs.

   The quantum de Finetti theorem shows that the essence of quantum-state tomography is not in revealing an "element of reality" but in deriving that various agents (who agree some minimal amount) can come to agreement in their ultimate quantum-state assignments. This is not the same thing as the stronger statement that "reality does not exist." It is simply that one need not go to the extreme of taking the "unknown quantum state" as being objectively real to make sense of the experimental practice of tomography.

   One is left with the feeling ... that perhaps this is the whole point to quantum mechanics. That is: Perhaps the missing ingredient for narrowing the structure of Bayesian probability down to the structure of quantum mechanics has been in front of us all along. It finds no better expression than in the taking account of the limitations the physical world poses to our ability to come to agreement.

I certainly believe there are some things within quantum mechanics that are beyond our subjective description. As in your example, Planck's constant could well be one of them. The dimensionality d of a Hilbert space is another one I feel fairly confident of. That number characterizes something intrinsic to a system. To that extent, it is not something that can be information-theoretic in origin.

Landahlism 4: *To press your point more forcefully in the future, you might consider rephrasing the special relativistic axioms themselves in a more information-theoretic light.*

But it doesn't seem to me that special relativity is overtly about Bayesian or information theoretic concerns in the way that quantum mechanics is. So I wouldn't want to express those axioms in a more information-theoretic light.

Landahlism 5: *Conclusion: quantum information theorists need to get out more!*

Yes! (And I speak for myself too.)

Landahlism 6: *At least I got a good groan from John Preskill when Alice flipped Bob a quarter for the "two bits" of classical communication she sent him.*

I got a similar groan from Mermin when he first read the end of Section 3. BTW, footnotes 8 and 9 are not typos; several people have asked me about that.

Landahlism 7: *In The Future section I talked about Gleason's theorem for POVMs and expressed my concerns about Emma's future psychological counseling given that you already are pressing her for a theory of measurement. None of us understood what the quote you ascribed to Hideo actually meant, which I suppose I should ask Hideo about. Do you understand it? It sure sounds amusing.*

The point is, Emma gets by without a theory of measurement, and we should all learn something from that. It is the people who think that knowledge acquisition, or better, belief acquisition, must arise from a detailed dynamical theory that are the problem.

The point of view taken here is that "detailed dynamical theories" are theories of *inference*, and therefore lie outside of the process of knowledge acquisition. This does not bar quantum theory from making contact with the REAL world—the world that was here long before man ever arose (see points about reality below)—it just means that one is not going to find it in the dynamics.

Landahlism 8: *In the Learning section, I went over your argument for quantum collapse being a kind of Bayesian conditioning. I understand the analogy you drew here, but I hardly believe this makes quantum collapse any more gentle of a process. That's because I don't believe Bayesian conditioning is "gentle." While it's true that one can express the classical process as "plucking" a term out of a sum over conditional probabilities, the change in probabilities can be quite dramatic: the change can be from nearly zero to one in a single step!*

Fair enough. Perhaps I overplayed the imagery. The point I really wanted to emphasize is that quantum collapse can be thought of as predominantly a refinement of one's knowledge.

Landahlism 9: *I also don't understand the meaning of the "mental readjustment" step in the quantum process. Is this just a change-of-basis for the description of the state? Is it something more or less than this? If that's all it is, then I really don't like this phrase "mental readjustment" at all.*

The track I'm on is that quantum state change is essentially Bayesian updating of knowledge, but with the proviso that the things we have to do to update our knowledge are (generally) not without effect on the world. This must be taken into account in some way. That the updating is Bayesian-like has a trace in quantum mechanics through Eqs. (57) and (58). That we still have to take into account our knowledge of our invasiveness, this has a trace in Eq. (59). That is the "mental readjustment"—i.e., taking into account what we know about our own invasiveness. (I agree, I should have found a better word for it.) When we know that our knowledge acquisition could not have physically affected the system it was concerned with, then we need do nothing whatsoever beyond Bayesian updating. Eq. (64) is an example of that.

Landahlism 10: *I finished the Learning section by sketching how "typical von Neumann entropy" as an uncertainty measure increases after every measurement. A question I had in your argument was why the integration is done over only von Neumann measurements. I looked over your original paper on this subject and didn't find an answer there either. Some mention is made of projective measurements being "maximally predictive," but I'm not totally convinced by this argument. I suspect that the true reason for restricting attention to these measurements is technical. A measure over POVMs doesn't exist, so one can't integrate over them. Wouldn't life be so much nicer if there were one! Is this the true reason for the von Neumann measurement restriction?*

You're suspicion is correct. My choice was no deeper than that.

Landahlism 11: *I spent most of my time discussing the Correlations section. I went in detail over your proof of what I called "Gleason's theorem for Classically Semilocalizable Operations (CSOs)" in deference to the terminology introduced in a recent paper by Beckman et al. Personally, I thought it was cool that the tensor product arose out of noncontextuality and the measurement model. John Preskill wasn't so impressed — he believes that the tensor product will arise out of any reasonable model of measurement which has the property of locality. (I.e where neither Alice's nor Bob's local actions can meaningfully impact the other.) He may be right, but for me that isn't the point. The point for me is that the proper way to view Gleason's theorem is as a machine. The input to the machine is the measurement model and the output of the machine is the state space structure and the probability law.*

Indeed you did get the point. Thanks. The point is, how much of the structure of quantum mechanics can we shove into the simple choice: "measurements = POVMs." How much of quantum mechanics is really independent of that choice? There has been a hell of a lot of work trying to reduce all of quantum mechanics to the assumption of unitarity. I'm trying to go the other way.

Landahlism 12: *What especially excites me about this point-of-view its potential impact on quantum field theory. The main point of the Beckman et al. paper is that causal measurements and localizable measurements are not one in the same. Wouldn't it be interesting to see what happens when we impose only causality on our measurement model and send it through the "Gleason machine?" What do you suppose the resultant structure of the state space would be?*

If I understand you correctly, Mermin in his talk in Montreal and Sweden has been wondering something very similar. In fact, he would like to see the quantum probability rule AND the tensor product rule arise out of the idea that measurement cannot give instantaneous signaling. He doesn't yet feel comfortable with my (Gleason's) noncontextuality assumption. Yours is a good question; I'll try to have a look at that paper.

Landahlism 13: *On the whole I portrayed your "party platform" as the statement that "Quantum states are states of knowledge about the consequences of future interventions."*

That statement, as it stands, is true.

Landahlism 14: *In particular, those consequences aren't consequences to reality, but rather consequences to states of knowledge about even further future interventions.*

That statement, as it stands, is not. (Do you not see the difference?!?!)

Landahlism 15: *In this worldview Bayesian agents don't work to align their predictions with an underlying reality.*

They would if they could, but they don't because they can't. Realizing this—it seems to me—is the first step to understanding what the quantum world is about.

Landahlism 16: *Instead they work to align their predictions with each other. It is as if reality in this picture is solely the agreement of predictions!*

I'd be interested to hear if you believe that this is a fair characterization of your party's platform. After reading this paper, I came to the conclusion that you didn't believe in reality at all. (Or at best I thought you believed reality = knowledge.) John Preskill tells me you believe otherwise, namely that there is a reality, which surprised me.

Yeah, you botch it pretty badly there. John is right. See my diatribe under C) above. But, let me also add to that:

1. Section 4, first paragraph.

   I have been watching my two year old daughter learn things at a fantastic rate, and though there have been untold numbers of lessons for her, there have also been a sprinkling for me. For instance, I am just starting to see her come to grips with the idea that there is a world independent of her desires. What strikes me is the contrast between this and the concomitant gain in confidence I see grow in her everyday that there are aspects of existence she actually *can* control. The two go hand in hand. She pushes on the world, and sometimes it gives in a way that she has learned to predict, and sometimes it pushes back in a way she has not foreseen (and may never be able to). If she could manipulate the world to the complete desires of her will, I am quite sure, there would be little difference between wake and dream.

This wispy little piece is the closest I've been able to come to giving substance what I call "the Paulian idea." See my "Notes on a ...", page vii. The world must have some impredictability about it, otherwise we would never be able to say we have seen any trustworthy trace of a reality.

Landahlism 17: *I'm curious to hear what you believe reality is.*

Me too. The idea is not well formed yet. Perhaps this accounts somewhat for people not getting my point that the first part of attempting to identify what is real in the quantum world is to identify what is subjective and governed by "laws of thought." We should do that because that's the easier part of the program. Contemplating what's left behind is when the real fun will begin.

People are too used to seeing gurus (like Deutsch or Mohrhoff) sit on high and declare what reality *is*. My goals are more modest, even if my method of advertisement is not. I don't have an answer yet; I just feel a direction. One should not confuse my method of attack with my answer.

If you were to push me real hard on this "what-you-believe-reality-is" business, I might be inclined to say, "Read Schopenhauer's *The World as Will and Representation*." But since I haven't read it myself, I can hardly expect you to do that! Anyway, as a very *provisional* answer, I might say it's something like the "will" (the quotes around that word are very important) that Schopenhauer attributes to every piece of the world, animate and inanimate alike. For want of a better term, I call it zing.

Landahlism 18: *As for the mechanical details of the paper itself, I enjoyed your refreshingly casual writing style. I found one of your section titles to be either exceedingly clever or merely a typographical error. Either way, I'm the only one out of a dozen people who noticed it, even after I pointed it out. I'm hoping that you intended the convey the clever interpretation.*

Caves, Bilodeau, and Schumacher also asked me if it is a typo. It is not: you can be relieved.

Landahlism 19: *If so, I suggest you correct the grammar on the section title to "Wither Entanglement!" to make the homonym less subtle without sacrificing any wit.*

Too late. But in any case, I wanted all the section headings to be questions, except the beginning and end ones.

Landahlism 20: *The only other typo that jumped out at me was on page 13: "shear difficulty" should read "sheer difficulty."*

Thanks, I hadn't noticed the difference before.

Landahlism 21 : *Once again, great paper. I'm psyched that it stirred up so much discussion in our group meeting.*

Me too!

Now, I said I would make some comments on Sumit's poll. But I'm too tired for that after all this writing. So I won't tell you what I think in any great detail at the present. I'll just cut and paste what Hans Primas thinks. It's below. I will say, however, that I don't see that there should be a qualitative distinction between my description of you (Andrew Landahl) and my description of the coffee maker sitting to my right. You are both physical systems embedded in this thing we call the world.

> From: H. Primas, "Beyond Baconian Quantum Physics," in *Kohti uutta todellisuuskäsitystä. Juhlakirja professori Laurikaisen 75-vuotispäivänä* (Towards a New Conception of Reality. Anniversary Publication to Professor Laurikainen's 75th Birthday), edited by U. Ketvel (Yliopistopaino, Helsinki, 1990), pp. 100–112.

The methodology of experimental scientific research and engineering science is to a large extent characterized by the regulative principles emphasized by Francis Bacon. It is a tacit assumption of all engineering sciences that nature can be *manipulated* and that the initial conditions required by experiments can be brought about by interventions of the world external to the object under investigation. That is, *we assume that the experimenter has a certain freedom of action which is not accounted for by first principles of physics*. Without this freedom of choice, experiments would be impossible. Man's free will implies the ability to carry out actions, it constitutes his essence as an actor. We act under the idea of freedom, but the topic under discussion is neither man's sense of personal freedom as a subjective experience, nor the question whether this idea could be an illusion or not, nor any questions of moral philosophy, but that *the framework of experimental science requires the freedom of action as a constitutive though tacit presupposition*.

The metaphysics of Baconian science is based on the confidence that only the past is factual, that we are able to change the present state of nature, and that nothing can be known about nature except what can be proved by *experiments*. Francis Bacon's motto *dissecare naturam* led to a preferred way of dividing the world into object and observing systems. An experiment is an *intervention* in nature, it requires artificially produced and deliberately controlled, reproducible conditions. In *experiments* in contradistinction to *observations* – one *prepares* systems in initial states, *controls* some of the variables, and finally *measures* a particular variable. The regulative principles of Baconian science require *power to create initial conditions*, stress *the facticity of the past* and *the probabilistic predictability of the future*, and reject *teleological considerations*.

10 July 2001, to Mermin, "Old McBleak's Ale House"

We are talking past each other.

But, my wording was careful enough to cover your reply (modulo the confusing parenthesis I put around the word quantum, for which I apologize). If a physical action associated with a POVM—by definition—only affects the system associated with the POVM's Hilbert space, then by definition that is all it affects. Standard quantum mechanics has that feature.

The issue is whether we should question the reasonableness of that. Or, indeed, as you would like, turn the tables and check whether the physical requirement of no-signaling gives rise to the standard probability rule full stop. Your question is a well-posed question, I do not deny that. But, as I view it, its motivation is a throw-back to the days when entanglement was thought to have some connection to the spooky ghosts of nonlocality.

I am torn. 75% of the time, I think your question is a regressive turn to the Popescu-Rohrlich-Aharonov-Shimony-Gisin "passion-at-a-distance" mentality. It seems to me acknowledging that as an interesting paradigm (even one to be ultimately shot down) is a wrong turn. But 25% of the time, I think, "Why not? It is a valid question, so answer it if you can."

Still no sympathy for me? (Probably not.) But, am I at least coherent?

18 July 2001, to Bub, "Horizons"

That is awful news about your visual problem and its possible causes. Please do keep me up-to-date on your health. For my own part I will cross my fingers and think of you often.

Don't worry at all about leaving me up in the air concerning a visit to Provence. Because of certain of our own medical issues, I've been lobbying my wife to postpone her European vacation until September or early October anyway. (So, you see, I would have to leave you up in the air right now too.) If it happens, it happens. The main thing is that it sounded like a good opportunity to pound out the similarities and distinctions between our points of view on quantum mechanics without being interrupted every three minutes.

I know I suggested I would write a longer letter soon, but I'm going to wimp out of it again for now. It would concern the main point of distinction I see between us (and also between myself and Pitowsky). Namely, A) that I view a large part of quantum mechanics as merely classical probability theory (which on my view may be an a priori "law of thought") PLUS an extra assumption narrowing down the characteristics of the phenomena to which we happen to be applying it to at the moment, while B) you are more tempted to view quantum mechanics as a *generalization* of classical probability theory (and with it information theory). I know that my view is not fully consistent yet, especially as I have always distrusted mathematical Platonism—which you pointed out to me I am getting oh so close to—but it still feels more right (to me, of course). Ben Schumacher, Ruediger Schack, and I had a long discussion on this (on a long walk) the day after the round table, and I'd like to record that too. Ben took a stance quite similar to yours, and maybe even Ruediger did too (despite his overwhelming Bayesianism). So, I may be the lonely guy out on this. And my view may be subject to change.

What I probably really need right now is more conversation than writing. So, I do hope I get to see you in an uninterrupted way soon. (By the way, would it be possible for me to get a copy of the talk you gave in Växjö? Could you copy that and mail it to me?) For now, let me post below parts of two notes I wrote Andrei Khrennikov. They touch on the discussion above, even if they are somewhat out of context here. The second note, in particular, struck me as a clean way of stating my position (in a way that I hadn't explored before). Maybe that'll help to zoom us in on the relevant issues.

There is some good news on the horizon. Gilles and I will be holding another foundations meeting in 2002. (Purely quantum info people, much like the original.) All the details aren't clear

yet, but we think we may be able to have desk space for people, it may be for an extended period—maybe a month—with a revolving set of participants, etc., etc. I hope you'll be able to join us (once I get the details to you). The main difficulty is that it may have to be in the fall (after the school semester starts), which will cause participation trouble for those with a teaching load.

18 July 2001, to Lee Smolin, "Page 270!!"

I'm excited to hear that you're reading some of my things. Especially since I've come across pages 270-272 of your book (paperback edition)! Indeed, there appears to be a significant overlap between some of our toy ideas. The one I'm speaking about is (in a technical way) the undercurrent of my paper "Quantum Foundations in the Light of Quantum Information." But, you can find broader-view statements of it on pages 156 and 190 of my samizdat. Also, you can see a trace of it in Mermin's foreword, page iii, last paragraph. Do you see the overlap that I do?

Anyway, I find this quite intriguing: Somehow, I had gotten the impression that you were a staunch many-worlder, and that our views of quantum mechanics might be diametrically opposite. I really apologize for my previous misreading (based on reading your book for an hour in a bookstore one day). I will read your book more carefully, and also look at the papers you recommended. [...]

You asked me, "What is the Paulian idea?" I wish I knew! More seriously, the best summaries I can give you are 1) the *conjunction* of two Pauli quotes on page vii of the samizdat, and 2) the wispy little piece I wrote in the first paragraph of Section 4 of my paper "Quantum Foundations in the Light ...". This is a very deep idea I think, and I don't know that I've ever seen it expressed anywhere except (very sketchily) in Pauli's writings. It is that, in a world where the experimental context cannot be deleted from the consequences it brings about, there must be a kind of randomness or impredictability. Else there would be no way to distinguish between wake and dream for any observer who makes use of such contexts. It is the ultimate impredictability of the consequences of our interactions with the world that gives us firm evidence that there is something beyond us. By this view, the world is not real because it can be mathematized completely, but because it cannot.

21 July 2001, to Landahl-Preskill, "The Reality of Wives"

This morning one of the local hospitals had a fund-raising flea market, and I picked up a copy of Martin Gardner's *The Whys of a Philosophical Scrivener* for $0.50. I haven't been able to put the thing down all day; it's quite good, and the beginning parts are especially relevant to my recent discussion with you.

This evening while sitting outside enjoying the end of the day, I couldn't help but read Kiki a cute little story from it. Gardner writes:

> When I was an undergraduate philosophy student at the University of Chicago I attended a seminar given by Bertrand Russell. Carnap, then a professor at Chicago, went to these sessions and often engaged Russell in spirited debates which I only partly comprehended. On one occasion they got into a tangled argument over whether science should assert, as an ontological thesis, the reality of a world behind the phaneron. [Phaneron was Peirce's term for the world of our experience, the phenomenal world.] Carnap struggled to keep the argument technical, but Russell slyly turned it into a discussion of whether their respective wives (Russell's new wife was knitting and smiling in a back-row seat) existed in some ontologically real sense or should be regarded as mere logical fictions based on regularities in their husbands' phaneron.

The next day I happened to be in the campus post office, where faculty members came to pick up mail. Professor Charles Hartshorne, a whimsical philosopher from whom I was then taking a stimulating course, walked in, recognized me, and stopped to chat.

"Did you attend the Russell seminar yesterday?" he asked. "I was unable to go."

"Yes," I said. "It was exciting. Russell tried to persuade Carnap that his wife existed, but Carnap wouldn't admit it."

Hartshorne laughed. Then, by a quirk of fate, in walked Carnap to get his mail. Hartshorne introduced us (it was the first time I had met Carnap; years later we would collaborate on a book); then, to my profound embarrassment, Hartshorne said: "Mr. Gardner tells me that yesterday Russell tried to convince you your wife existed, but you wouldn't admit it."

Carnap did not smile. He glowered down at me and said, "But that was not the point at all."

I followed that by saying, "You know some of my friends are afraid that I don't believe in reality. So there, you're just a figment of my imagination!" She reacted in shock. "Well, I know that can't be true," she said. "Clearly you'd make some changes!"

22 July 2001, to Mermin, "Noncontextual Sundays"

I know you're busy, but I'm going to try again. (Don't feel the need to write back until you get some time.) The issue is still noncontextuality in the Gleason-like theorems: Is it a natural assumption or not?

Here was the best answer I gave you before, but now I'm going to try to improve on it.

> **Merminition 4:** *The question you're evading is what it means for one and the same positive operator E to appear in many different POVMS.*
>
> No, I don't think I'm evading it. It means that those various interventions or ways of gathering data—those POVMs—physically diverse though they may be—all lead to at least one common possibility for what my knowledge can be updated to (modulo the unitary readjustment).

The point I'm going to try to make is that not only am I finding noncontextuality a natural assumption, but actually it may be the most *basic* assumption of the whole game. (I.e., it may even be prior to the notion that measurements correspond to POVMs.) The idea is captured above, but—I can see now—it is in too idiosyncratic of a language to convince you easily.

Here's the new shot at it (emphasizing a slightly different aspect than previously).

1) Here's the scenario. Forget about quantum mechanics for the moment. Let me take a system S and imagine acting on it with one of two machines, M and N. For the case of machine M, let us label the possible consequences of that action $\{m_1, m_2, \ldots\}$. For the case of machine N, let us label them $\{n_1, n_2, \ldots\}$.

2) If we are good Bayesians, nothing will stop us from using all the information available to us to ascribe probabilities to the consequences of these two potential actions. Thus we naturally have lying around two probability distributions, $p_M(m_k)$ and $p_N(n_k)$.

That's well and good, but it's hardly a physical theory yet. We need more. So, let us suppose the labels m_k and n_k are at least drawn from the same master set (possibly even a set with further

structure, like a vector space or something). But then we must ask, under what conditions should we identify two particular labels m_i and n_j with the same element in the master set?

There's really only one thing lying around to do it with, and that's the probability assignments. If $p_M(m_i) \neq p_N(n_j)$, then surely we would not imagine identifying m_i with n_j. If, on the other hand, $p_M(m_i) = p_N(n_j)$ *regardless* of the initial state of knowledge about S, then we might think there's some warrant for it.

And that's the whole story of noncontextuality. It is nothing more than: The consequences (m_i and n_j) of our disparate actions (M and N) should be labeled the same when we would bet the same on them in all possible circumstances (i.e., regardless of our initial knowledge of S).

By this point of view, noncontextuality is a tautology—it is built in from the start. Asking why we have it is a waste of time. Where we do have a freedom is in asking why we make one particular choice of a master set over another. Why should the m_i's be drawn from the set of "effects" (i.e., the positive operators smaller than the identity on some Hilbert space)? Recall the problem on page 86 of the samizdat. Not all choices of the master set are equally interesting once we've settled on noncontextuality for the probability assignments.

You see, I really never do dismiss anything you say! Now I'm off to buy a new BBQ grill. (My family is tugging on me, and refused to let me try to hone this letter. But I really hope by this point it does make some sense.)

23 July 2001, to Summhammer, "Law without Law"

I very much enjoyed reading your letter to Carl Caves titled "promoting the Bayesian view." It was quite thoughtful, and makes me regret not having talked to you more while we were in Sweden.

Please allow me to ask a couple of questions based on what you wrote.

Summhammerism 1: *Here I think that, even if there is absolutely no order in the physical world, it will exhibit statistical order to a rational observer. The existence of rationality is to be taken as outside the physical word. It is a transcendental fact. (The term "law of thought" in your paper circles around the same thing.) Analysis of brain functions and molecules explains nothing, because that analysis is done by means of rationality.*

There have been times in my life when I have been very attracted to ideas like this. In particular, right now might be one of them (though I have a history of going up and going down). My first exposure to the idea came from John Wheeler's writings on what he termed "law without law." In fact presently, I'm putting together a large compendium of quotes and citations titled "The Activating Observer: Resource Material for a Paulian-Wheelerish Conception of Nature." At the moment, it consists of 423 annotated citations, taking up 96 pages of print. The manuscript is far from complete, but will eventually be submitted to *Studies in History and Philosophy of Modern Physics*.

I wonder if you have any suggestions for things I should include in it (based on your passage above). If so, please give me as complete of references as possible. If I'm not mistaken, I detect a Kantian tinge in your thought: that's an area I haven't explored too deeply in my compendium.

Summhammerism 2: *p.21: ... the Bayesian interpretation places actualization outside its provenance ...*

I agree. Defenders of interpretations which claim to handle actualization should read selected articles on the mind-body problem, written over the last three thousand years.

Would you mind expanding on this, and also what are some of those selected articles? Can you provide references?

Anyway, it was very nice meeting you for the first time.

23 July 2001, to Caves, "A Nonbayesian Bayesian?"

I enjoyed reading your dialogue with Summhammer.

I have one question of my own.

Cavesism 1: *I do believe that natural selection only works in a world with at least a statistical order, which leads to the quasi-determinism of the macroscopic world. That we and other creatures are exquisitely attuned to this order, to the point of often finding it where it's not really there, is not surprising. The costs of finding order where there is only chaos must be less than the cost of failing to notice and take advantage of order when it is there.*

What on earth do you mean by this? Similarly when you write, "statistical order is the first element of the Bayesian reality," in Section 7 of your *Resource Material*?

Your phrase "world with statistical order" seems to teeter awfully close to an objectivist notion of probability. And it frightens me, of course, having fully made a conversion now.

23 July 2001, to Caves, "The Principal PrincipleS"

It dawned on me that I should chide you on your discussion on page 21, starting with "Two further comments on Hamiltonians ...." You left out the all-important part about how to connect the two notions of Hamiltonian, the subjective—or effective, as you call it—with the objective. It must be a Cavesian version of the principal principle: When the objective determinant of the time evolution of one's subjective states of knowledge is known, then the subjective determinant of one's subjective states of knowledge should coincide with it.

Another typo btw (I think). Page 19: "If you have maximal information about a quantum system and you want to retain it, you must know the system is Hamiltonian." Don't you mean "system's Hamiltonian"?

24 July 2001, to myself, "Feynman Quotes"

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generation of creatures, what statement would contain the most information in the fewest words? I believe it is the atomic hypothesis (or the atomic fact) that all things are made of atoms—little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into on another.
>
> Everything is made of atoms. That is the key hypothesis.

30 July 2001, to Renes, "Britannica"

This was the only passage in the *Encyclopedia Britannica* that I could find about Nietzsche that even remotely resembled quantum mechanics. In general, he looks like a lot to wade through for little return.

> Perspectivism is a concept which holds that knowledge is always perspectival, that there are no immaculate perceptions, and that knowledge from no point of view is as incoherent a notion as seeing from no particular vantage point. Perspectivism also denies the possibility of an all-inclusive perspective, which could contain all others and, hence, make reality available as it is in itself. The concept of such an all-inclusive perspective is as incoherent as the concept of seeing an object from every possible vantage point simultaneously.
>
> Nietzsche's perspectivism has sometimes been mistakenly identified with relativism and skepticism. Nonetheless, it raises the question of how one is to understand Nietzsche's own theses, for example, that the dominant values of the common heritage have been underwritten by an ascetic ideal. Is this thesis true absolutely or only from a certain perspective? It may also be asked whether perspectivism can be asserted consistently without self-contradiction, since perspectivism must presumably be true in an absolute, that is a nonperspectival sense. Concerns such as these have generated much fruitful Nietzsche commentary as well as useful work in the theory of knowledge.

02 August 2001, to Schumacher, "Re: The Montréal Commune – Ditto"

I've marked you down in my spreadsheet, and we'll let you know what's up in a couple of weeks. I like your result! (which can't be questioned)

> based on two propositions:
> (1) "Information" resides in the relation between systems, and (2)

No subject without an object. No object without a subject. No information without both. I like that—it seems like a good track for ontologizing information, to the extent that it can be. The only thing that scares me is your secret desire to reify the quantum state—namely by translating (1) into a statement about bipartite quantum states (which seem to me to have no other good interpretation than information to begin with)! It takes information to get information off the ground?

07 August 2001, to Brun-Finkelstein-Mermin, "Knowledge, Only Knowledge"

Below is a note I started composing last Friday—but then had to leave for a long weekend for my wife's birthday—and only finished up today. In the mean time, Todd and Jerry have skirted very close to the point I wanted to make. So, the note is not quite as relevant as it might have been, but maybe some of it is still worth contemplating.

Allow me to start off in a fanciful way (like usual) with a couple of quotes:

> The subjectivist, operationalist viewpoint has led us to the conclusion that, if we aspire to quantitative coherence, individual degrees of belief, expressed as probabilities, are inescapably the starting point for descriptions of uncertainty. There can be no theories without theoreticians; no learning without learners; in general, no science without scientists. It follows that learning processes, whatever their particular concerns

> and fashions at any given point in time, are necessarily reasoning processes which take place in the minds of individuals. To be sure, the object of attention and interest may well be an assumed external, objective reality: but the actuality of the learning process consists in the evolution of individual, subjective beliefs about that reality. However, it is important to emphasize, as in our earlier discussion in Section 2.8, that the primitive and fundamental notions of *individual* preference and belief will typically provide the starting point for *interpersonal* communication and reporting processes. In what follows, both here, and more particularly in Chapter 5, we shall therefore often be concerned to identify and examine features of the individual learning process which relate to interpersonal issues, such as the conditions under which an approximate consensus of beliefs might occur in a population of individuals. — pp. 165–166, Bernardo and Smith, *Bayesian Theory*
>
> What is the nature and scope of Bayesian Statistics within this spectrum of activity? Bayesian Statistics offers a rationalist theory of personalistic beliefs in contexts of uncertainty, with the central aim of characterising how an individual should act in order to avoid certain kinds of undesirable behavioural inconsistencies. The theory establishes that expected utility maximization provides the basis for rational decision making and that Bayes' theorem provides the key to the ways in which beliefs should fit together in the light of changing evidence. The goal, in effect, is to establish rules and procedures for individuals concerned with disciplined uncertainty accounting. The theory is not descriptive, in the sense of claiming to model actual behaviour. Rather, it is prescriptive, in the sense of saying "if you wish to avoid the possibility of these undesirable consequences you must act in the following way." — p. 4, Bernardo and Smith, *Bayesian Theory*

Thanks again to all of you for letting me look in on your interesting emails! I've learned a lot from this exchange. Last night I was up between 1:30 and 5:00 reading them all one more time (and David's original paper too), and thinking much harder than I had before on these issues. So, now I hope to say some things in that regard, but not inane things (as I had done a couple of days ago).

The main thing that started striking me more deeply last night is that now it is very obvious that you *have* an answer, but—much more than ever before—I don't really understand what the question ought to be. In a sense, I'm only coming across the same troubles (Samizdat, p. 236) I've had ever since David first wrote me on this issue, and it relates to the comment I made after his talk in Växjö (in case he remembers). Let me try to explain.

David, in his paper, quotes Peierls as saying,

Merminition 5 : *In my view the most fundamental statement of quantum mechanics is that the wavefunction, or, more generally the density matrix, represents our knowledge of the system we are trying to describe. . . . [Yet, density matrices] may differ, as the nature and amount of knowledge may differ. People may have observed the system by different methods, with more or less accuracy; they may have seen part of the results of another physicist. However, there are limitations to the extent to which their knowledge may differ.*

And David himself says,

Merminition 6 : *I have the feeling that if quantum mechanics is really about knowledge and only knowledge, then there ought to be further elementary constraints on the possible density matrices*

describing one and the same physical system that are stronger than the very weak second condition of Peierls, but not as strong as his overly restrictive first condition.

What is being called for here—perhaps unintentionally—is a way to think about quantum mechanics from the bottom up (as Caves, Schack, and I might like it, in a Bayesian way), rather than from the top down (as Everett, Deutsch, and Bennett might like it). That is, one should view quantum mechanics as a conduit for stitching our individual pictures/thoughts/beliefs into a pastiche we ultimately call "the world." This contrasts with imagining that we have miraculously grasped the ultimate reality (the universal wavefunction, say), and can somehow see our individual points of view as being derived back out of that.

But if this is the case, then I cannot understand Peierls' command, David's quest, or the *answer* all three of you ended up coming up with:

> Two density operators ρ_a and ρ_b can describe the same system if and only if the support of ρ_a has a nontrivial intersection with the support of ρ_b.

This theorem is certainly consistent with a top-down view of the Everett sort—that, to be explained below, is really is what it seems to me to demonstrate—but it is not consistent with the bottom-up view. For, from the bottom-up view, there should not—and more importantly, there cannot—be any constraints whatsoever on what an agent can believe. I have every right to be as wrong-headed as I want to be with respect to you: The density operator belongs to me, not to you, and not to the system. I have every right to say inane things and make inane predictions—I do it all the time. What I should not do, however, if I want to remain rational, is refuse to listen to you when you point out my inanities or refuse to listen to the detector clicks that contradict my previous predictions. From the Bayesian view, it is the process of updating and the general structure of beliefs that is constrained by rationality—i.e., by the physical world, or the Platonic ideal, depending upon your orientation. It is not the actual beliefs themselves.

This point, perhaps more than any other, is why I (and Caves and Schack) should adopt the word "belief"—rather than "knowledge" or "information"—for describing the operational significance of quantum states. (I will try to be more consistent in the future, but that is really an aside as far as this note is concerned.)

Of course, your theorem is a theorem—or at least I can see nothing wrong with it—the issue here is how I, with my little Bayesian mind, can put it into a context I am more happy with. At first, I was the most pleased with Jerry's way of motivating it: From that point of view, what the theorem seems to express is simply the conditions under which a third agent Carol can consistently incorporate Alice and Bob's disparate beliefs into her own belief system.

But from the Bayesian view, why should we care about a Carol at all? What if there's no Carol to be found? What if neither Alice nor Bob ever intend to share their thoughts about this poor physical system with anyone else? To say that there is always a Carol about, or that there ought to be one, is to come dangerously close to endorsing the Everettian (or, for that matter, Bishop Berkeleyian) program. This, of course, may not bother Jerry or Todd—I'm not completely sure about their foundational dispositions—but it does bother me, and I suspect it might bother David, with his newfound deconstructionist tendencies.

Thus, it now seems to me that Todd's original way of posing the issue may be the safer way after all. BUT that is not because it gives us *the* answer, but instead *an* answer. (I.e., there are should *only* be sufficient conditions, rather than necessary ones.) Let me let Bernardo and Smith speak again (and again):

> [T]here is an interesting sense, even from our standpoint, in which the parametric model and the prior can be seen as having different roles. Instead of viewing these roles as corresponding to an objective/ subjective dichotomy, we view them in terms of an intersubjective/ subjective dichotomy. To this end, consider a *group* of Bayesians, all concerned with their belief distributions for the same sequence of observables. In the absence of any general agreement over assumptions of symmetry, invariance or sufficiency, the individuals are each simply left with their own subjective assessments. However, given some set of common assumptions, the results of this chapter imply that the entire group will structure their beliefs using some common form of mixture representation. Within the mixture, the parametric forms adopted will be the same (the *intersubjective* component), while the priors for the parameter will differ from individual to individual (the *subjective* component). Such intersubjective agreement clearly facilitates communication within the group and reduces areas of potential disagreement to just that of different prior judgements for the parameter. As we shall see in Chapter 5, judgements about the parameter will tend more towards a consensus as more data are acquired, so that such a group of Bayesians may eventually come to share very similar beliefs, even if their initial judgements about the parameter were markedly different. We emphasize again, however, that the key element here is intersubjective agreement or consensus. We can find no real role for the idea of objectivity except, perhaps, as a possibly convenient, but potentially misleading, "shorthand" for intersubjective communality of beliefs. — pp. 236–237, Bernardo and Smith, *Bayesian Theory*

> In the approach we have adopted, the fundamental notion of a model is that of a predictive probability specification of observables. However, the forms of representation theorems we have been discussing provide, in typical cases, a basis for separating out, if required, two components; the parametric model, and the belief model for the parameters. Indeed, we have drawn attention in Section 4.8.2 to the fact that shared structural belief assumptions among a group of individuals can imply the adoption of a common form of parametric model, while allowing the belief models for the parameters to vary from individual to individual. One might go further and argue that without some element of agreement of this kind there would be great difficulty in obtaining any meaningful form of scientific discussion or possible consensus. — p. 237, Bernardo and Smith, *Bayesian Theory*

You'll find something similar to this theme infused throughout my paper `quant-ph/0106166`, "Quantum Foundations in the Light of Quantum Information." In the present context the question is, under what conditions (on the density operators themselves), can Alice and Bob move toward consensus in their future density operator assignments for a system? You guys have answered this for the *case* when there are some extra systems beside the one of interest (all noninteracting) lying around for Alice and Bob to make measurements on. If there is no overlap between the initial supports, then no measurement-at-a-distance, which can only have the effect of refining a density operator, can get Alice and Bob any closer to agreement.

What about necessity? As I've tried to say, there can be no requirement of necessity from the Bayesian view I'd like to see prevail. Instead the question is *always* to identify those situations where agents (with disparate beliefs) can move toward consensus be it by indirect measurements, direct ones, or even by communication with further members of the community.

A good and indicative example comes from the quantum de Finetti theorem along with the points made by `quant-ph/0008113` (Schack, Brun, and Caves). From it, we have a natural case

where two agents can start out with distinct density operator assignments on a large collection of systems, but through updating via some commonly viewed measurement outcomes, they can move toward complete agreement in their estimates of the outcomes of all future measurements. The only thing the agents need to walk into the room with is this much INITIAL AGREEMENT: (1) that the systems in the collection are exchangeable, and (2) that the parametric form given by the de Finetti theorem has full support (in the sense defined there). Without that initial agreement, the techniques of quantum-state tomography would lead to no final agreement at all. No one can require that two observers in Mike Raymer's lab must walk into it with such an initial agreement, but if they happen to, there will be a reward at the end of the day.

That's my spiel. How would I modify David's quest in light of all that I've just said? Here's my shot:

> I have the feeling that if quantum mechanics is really about knowledge and only knowledge—or better, belief convergence and *only* belief convergence—then FOR ANY GIVEN METHOD OF GATHERING INFORMATION, there *should* be a way to ferret out of quantum mechanics the necessary and sufficient conditions on two observers' initial state assignments, so that the gathered information leaves them in a better agreement than they started out with.

Let me give you an example of where such an exercise can go. This is a question I posed to Ruediger while we were in Växjö, but also it is a generalization of Todd's considerations (now letting Alice and Bob's measurements being disturbing ones).

Let us agree on a distance measure on density operators. A convenient one (and my favorite) is

$$d(\rho,\sigma) = \mathrm{tr}|\rho-\sigma|,$$

but that's not so important for the considerations here. (Well, it may be in the long-run, but I don't want to let such a technicality detract from posing the question.)

Suppose Alice walks into a room and says to herself that a system is described by ρ, while Bob says to himself that it is described by σ. We can gauge their amount of initial consensus by $d(\rho,\sigma)$. Now, suppose Carol (or Alice or Bob for that matter) performs a measurement on the given system whose action—i.e., whose associated completely positive map—is to take any initial state τ to

$$\tau \longrightarrow \tau_b = \frac{1}{\mathrm{tr}(\tau A_b^\dagger A_b)} A_b \tau A_b^\dagger$$

depending upon the particular outcome b. Here, of course,

$$\sum_b A_b^\dagger A_b = I.$$

(This is nothing more than the general form of an "efficient" measurement as defined in my QFILQI paper cited above.) Let us allow Alice and Bob to be privy to this map, and indeed to the actual outcome b.

Thus, the final consensus of Alice and Bob will be gauged by $d(\rho_b,\sigma_b)$. The question is: When can Alice and Bob expect to be in better agreement after the measurement than before? That is, as far as Alice is concerned, she will *expect* their final distance to be

$$D_A = \sum \mathrm{tr}(\rho A_b^\dagger A_b) d(\rho_b,\sigma_b).$$

As far as Bob is concerned, he will *expect* their final distance to be

$$D_B = \sum \mathrm{tr}(\sigma A_b^\dagger A_b) d(\rho_b,\sigma_b).$$

For a given set of A_b's, what are the necessary and sufficient conditions on ρ and σ so that

$$D_A \leq d(\rho, \sigma)$$

and

$$D_B \leq d(\rho, \sigma) \ ?$$

Suppose we can answer these questions. Then we will be able to identify the minimal fact Alice and Bob must reveal to each other (even without explicitly revealing their full beliefs captured by ρ and σ) so that they can expect to walk out of the room in better agreement ... even if they still can't say with certainty what each other now believes.

Beliefs, only beliefs. But sometimes we can say something about their convergence. And you guys have provided an example.

07 August 2001, to Brun-Finkelstein-Mermin, "Amendment"

I just reread the thing after sending it. Let me be more careful before one of you accuses me of being a flaming positivist again. (I suspect Todd, in particular, will be inclined to do so, and I want to fend that off before it happens.) In my closing sentence, I wrote:

> Beliefs, only beliefs. But sometimes we can say something about their convergence. And you guys have provided an example.

Let me temper that to:

> Beliefs, only beliefs. But sometimes we can say something about their convergence, as they are steered by our interactions with the world external to us. And you guys have provided an example.

Now, I can sleep safely ...
Good wishes to all!

07 August 2001, to Finkelstein, "The First Amendment"

Thanks for the note. But nothing changes for me (yet, at least). Sorry for my overemphasis on Carol.

Finkelsteinism 5: *If that can be agreed to, then the rest is essentially just algebra. $S[rho_a]$ is the orthogonal complement of the zero-eigensubspace of rho_a, which is the set of $|\phi\rangle$ such that Alice knows with certainty that a measurement of $|\phi\rangle\langle\phi|$ must yield the result zero. So the support of the updated version of ρ_a must be a subset of the support of the original ρ_a. Likewise, Alice can see that the support of the updated rho must be a subset of the support of ρ_b. Etc.*

The issue for me is not whether Alice *might* be willing to incorporate Bob's beliefs/knowledge into her own knowledge base—that, I'll grant her, in which case I agree that everything you say is true. Instead it is whether she *must* be willing to do it. I think you are (tacitly) trying to get me to agree that she *must* be willing to accept Bob's quantum state as extra, valid information. That's something I can't do. I reserve the right for Alice to think that Bob's quantum state is complete nonsense, something that she would never want to incorporate into her own knowledge base. (It's part of being an American.)

To say that two observers *must* be willing to incorporate their separate states of knowledge into a single state is—I think—to tacitly (there's that word again) accept an Everettian kind of view. For it never allows that quantum states are states of belief in the normal sense. Instead it makes them more like "objective points of view" (relative states) that *must* be derivable from a larger, more encompassing picture. Why else would the states have to be "consistent" with each other (for instance, in your and Todd's sense)?

Is that helping any to get my wacky point of view across?

07 August 2001, to Waskan, "Kiki, James, and Dewey"

The last month or so, I've been logging quite a few hours in the philosophical world. I read about 100 pages from this little book: A. Schopenhauer, *The Philosophy of Schopenhauer*, edited, with an introduction by I. Edman, (Modern Library, New York, 1928).

It's a small collection of pieces from Schopenhauer's big masterwork. I read about 50 pages on the world as representation, and about 50 on the world as will. Even that little was not easy for me. It's fanciful stuff, but maybe I extracted an idea or two that I like. The main one is simply the idea of a dichotomy between what things look like from the inside of any phenomenon (when there is a view from that perspective)—Schopenhauer called it will—and what things look like from the outside of the phenomenon (when there is a view from that perspective)—Schopenhauer called it idea. Or, maybe I should have more safely said "is" rather than "look like." But, in any case, that distinction (and a strict separation between the concepts) strikes me as useful or at least worth contemplating. Just about all the rest, though, I probably wouldn't be able to accept: the strict Kantian categories, the principle of sufficient reason, etc., etc.

On the other hand, I have gotten *completely* carried away with William James and John Dewey. Here's how I put it to my friend Carl Caves the other day:

> Today I focused on rounding up some more William James, John Dewey, Percy Bridgman material. I think James is taking me over like a new lover. I had read a little bit of him before, but I think I was more impressed with his writing style than anything. But I was drawn back to him by accident, after reading Martin Gardner's *Whys of a Philosophical Scrivener*. Gardner devoted a lot of time knocking down James' theory of truth, because it is just so much easier to accept an underlying reality that signifies whether a proposition is true or false, rather than saying that the knowing agent is involved in eliciting the very proposition itself (along with its truth value). And something clicked! I could see that what James was talking about might as well have been a debate about quantum mechanics. He was saying everything in just the right way. (Let me translate that: he was saying things in a way similar to the way I did in my NATO "appassionato.") And things have only gotten better since.

And indeed, they have only gotten better since! Since coming to Munich, I have not been able to put James and Dewey down (when I'm not writing emails trying to translate their ideas into the quantum mechanical context, in particular for a technical problem Caves, Schack and I are disagreeing violently on). I read James' *Pragmatism, and four essays from the Meaning of Truth*, and now I'm about halfway through *John Dewey: The Essential Writings*. I'm moved by this stuff like nothing else I've ever read.

You can't tell me philosophers don't have the good life!

08 August 2001, to Finkelstein, "Cross Entropy Min"

Finkelsteinism 6: *Which brings up a slightly different question. Suppose, again, that Alice describes a system by ρ_a, but this time let's say for simplicity that she considers what would happen if Bob were to tell her that his ρ_b corresponded to a pure state $|\phi\rangle$; then, given her knowledge of the system (ie, given ρ_a) what could Alice say about what $|\phi\rangle$ might be?*

It follows from what we have been saying that Alice knows that for $|\phi\rangle$ to be possible, it must be in $S[\rho_a]$, and that (in finite dimensions, at least) she cannot, based on her own knowledge, rule out any $|\phi\rangle$ that is in $S[\rho_a]$. But can Alice say any more than that? Would it make sense for Alice to put a probability distribution on the possible $|\phi\rangle$ that Bob might announce to her? (It would have to be a probability density; eg, if ρ_a were a multiple of the identity, then Alice would surely judge all states to be equally-likely.) If that did make sense, it would mean that Alice would be constructing a particular (continuous) ensemble representation of ρ_a. What could that be?

As I recall, this is quite similar to the classical problem ("principle of minimum cross-entropy") explored in the two references below. But I'm not going to have a chance to refresh my memory for a while. If you've got the time, you might see if it's relevant.

1. John E. Shore and Rodney W. Johnson, "Axiomatic Derivation of the Principle of Maximum Entropy and the Principle of Minimum Cross-Entropy," IEEE Transactions on Information Theory **IT-26**(1), 26–37 (1980).

2. John E. Shore and Rodney W. Johnson, "Properties of Cross-Entropy Minimization," IEEE Transactions on Information Theory **IT-27**(4), 472–482 (1981).

08 August 2001, to Caves, "The First Eye"

I am just about to get down to some serious (political) scheming to do with our Bayesian program: I'll let you know what I'm talking about if it turns out to be successful. First though, I want to take a moment to tell you about a point of similarity between our (both far-from-completely worked out) flavors of quantum ontology. This one just struck me a few days ago.

Some time ago, I tried to explain to you what I was hoping for for [sic] an ontology behind quantum mechanics. (See Samizdat, pp. 127–129.) I said it would have something to do with the information-disturbance tradeoff in quantum eavesdropping. You replied that, try as you might, you could see no ontological content in such a statement. I think what was troubling you was that the information-disturbance relations (as I am thinking of them), by their very nature, require explicit reference to the *subjective* points-of-view/opinions/beliefs (i.e., quantum states) of various *agents*.

I, on the other hand, have no problem with that. For the way to think about it is that the world (independent of our existence) has latent within it a property that simply has no way of being properly expressed without inserting information-manipulating agents into the picture. (Or, at the very least, that this anthropocentric way of stating things may be our first firm handle for getting at a better, more objective-sounding, formulation of the latent property.) Taking this tack does not mean, of course, that there is no world independent of human existence, and that is my point. It just means that we may sometimes have to take into account our (presumably contingent) existence for expressing some of the world's properties.

Here is something, however, that you should think about in connection to your "world = Hamiltonian" hopes. Let me put the ball back into your court. Can you explain to me the role of

Hamiltonians in your ontology in a way that does not make use—even tacitly—of the concept of a (subjective) quantum state? What is it that Hamiltonians do if their primary role is not in evolving (subjective, agent-required) quantum states? In the classical world, one could give an answer to this question by saying, "They evolve the positions of phase-space points." Such a statement makes no use of the concept of information-bearing agents for its formulation. (I view it as just nitpicking to argue whether the points or their trajectories (from which we can derive the Hamiltonian) are the more primary of the entities. Or whether they are equally primary.) But in the quantum case, I haven't yet seen what you can say if I take away the linguistic tool of "the quantum state" from your explanatory repertoire. What can you say?

I'll put a few passages of Schopenhauer below to inspire you. [From: A. Schopenhauer, *The Philosophy of Schopenhauer*, edited, with an introduction by I. Edman, (Modern Library, New York, 1928).]

> "No object without a subject," is the principle which renders all materialism for ever impossible. Suns and planets without an eye that sees them, and an understanding that knows them, may indeed be spoken of in words, but for the idea, these words are absolutely meaningless. On the other hand, the law of causality and the treatment and investigation of nature which is based upon it, lead us necessarily to the conclusion that, in time, each more highly organised state of matter has succeeded a cruder state: so that the lower animals existed before men, fishes before land animals, plants before fishes, and the unorganised before all that is organised; that, consequently, the original mass had to pass through a long series of changes before the first eye could be opened. And yet, the existence of this whole world remains ever dependent upon the first eye that opened, even if it were that of an insect. For such an eye is necessary condition of the possibility of knowledge, and the whole world exists only in and for knowledge, and without it is not even thinkable. The world is entirely idea, and as such demands the knowing subject as the supporter of its existence. This long course of time itself, filled with innumerable changes, through which matter rose from form to form till at last the first percipient creature appeared—this whole time itself is only thinkable in the identity of a consciousness whose succession of ideas, whose form of knowing it is, and apart from which, it loses all meaning and is nothing at all. Thus we see, on the one hand, the existence of the whole world necessarily dependent upon the first conscious being, however undeveloped it may be; on the other hand, this conscious being just as necessarily entirely dependent upon a long chain of causes and effects which have preceded it, and in which it itself appears as a small link. These two contradictory points of view, to each of which we are led with the same necessity, we might again call an *antinomy* in our faculty of knowledge, and set it up as the counterpart of that which we found in the first extreme of natural science. The objective world, the world as idea, is not the only side of the world, but merely its outward side; and it has an entirely different side—the side of its inmost nature—its kernel—the thing-in-itself. This we shall consider in the second book, calling it after the most immediate of its objective manifestations—will. But the world as idea, with which alone we are here concerned, only appears with the opening of the first eye. Without this medium of knowledge it cannot be, and therefore it was not before it. But without that eye, that is to say, outside of knowledge, there was also no before, no time. Thus time has no beginning, but all beginning is in time. Since, however, it is the most universal form of the knowable, in which all phenomena are united together through causality, time, with its infinity of past and future, is present in the beginning of knowledge. The

> phenomenon which fills the first present must at once be known as causally bound up with and dependent upon a sequence of phenomena which stretches infinitely into the past, and this past itself is just as truly conditioned by this first present, as conversely the present is by the past. Accordingly the past out of which the first present arises, is, like it, dependent upon the knowing subject, without which it is nothing. It necessarily happens, however, that this first present does not manifest itself as the first, that is, as having no past for its parent, but as being the beginning of time. It manifests itself rather as the consequence of the past, according to the principle of existence in time. In the same way, the phenomena which fill this first present appear as the effects of earlier phenomena which filled the past, in accordance with the law of causality. Those who like mythological interpretations may take the birth of Kronos ($\chi\varrho o\nu o\varsigma$), the youngest of the Titans, as a symbol of the moment here referred to at which time appears, though indeed it has no beginning; for with him, since he ate his father, the crude productions of heaven and earth cease, and the races of gods and men appear upon the scene.

10 August 2001, to Brun, "The Fifth Amendment"

Thanks for the note, which I thoroughly enjoyed. You hit a lot of nails on the head with it. Let me try to expand on some of the points that—I believe—show that at least you and I are coming to a little consensus. (As for Jerry and David, I will put them in a superposition for the time being, and see how this this interaction shakes things up.)

Brunism 1: *That doesn't mean that I completely disagree with you, Chris, but I think you are making a point which is pretty far from the spirit of this problem.*

This is the point. What I am trying to get straight is: *What* is the spirit of this problem? You write,

Brunism 2: *We have been describing a consistency criterion.* If *one wishes to combine two state descriptions of a single system into a* single *state description, the criterion tells one* when *it is consistent to do so (i.e., when the two descriptions are not actually contradictory).*

I agree that nobody is holding a gun to Alice's head and forcing her to incorporate Bob's information.

Putting it like that, I can certainly accept the proposition. I want to emphasize that. The attractive feature for me is that it is built on a conditional at the outset.

But you speak of *the* spirit of *this* problem. How is your statement to be reconciled with the tone of David's `quant-ph`? In particular, say, David's quote of Peierls (which he takes as his guiding light):

Merminition 7: *In my view the most fundamental statement of quantum mechanics is that the wavefunction, or, more generally the density matrix, represents our knowledge of the system we are trying to describe. ... [Yet, density matrices] may differ, as the nature and amount of knowledge may differ. People may have observed the system by different methods, with more or less accuracy; they may have seen part of the results of another physicist. However, there are limitations to the extent to which their knowledge may differ.*

Peierls does not use the qualification that you did—nor do any of you three in many of the emails I have seen—and that is what bothers me.

What I see in Peierls version of the spirit is that two quantum states cannot co-exist (even in a Platonic sense) unless they are consistent (in one manner or another, yet to be fleshed out). It is as if the universe has these little properties floating about, called quantum states, that MUST be consistent in the BFM sense (or some other sense). I will agree that that might be fine from an Everett-kind of point of view. But if one insists on consistency (as one should with ontological, physical properties), then—it seems to me—one breaks away from the desire to give the quantum state a *purely* epistemological role ... which, as I understood it, was the goal of Peierls and Mermin (though I am willing to accept that it may not be the goal for you and Jerry).

Brunism 3: *The point of consistency is to determine* if *two points of view can be combined into a single description, not to require that they must be.*

This I also agree with: It was meant to be the whole point of my note "Knowledge, Only Knowledge." My only point of disagreement, was that it was seeming to me that the tendency in David (and Jerry?) was the assumption that the "if" must be satisfied—in some sense—in the "real" physical world.

I just don't know how to say this more clearly. I think it is a valid worry about the intent, the very definition of what the problem is about. I don't think I'm being subtle: that is certainly the last thing I want to be.

Brunism 4: *I will make one additional comment. I think that in science there is usually a* tacit *(to use your word) assumption that the separate states of knowledge of different observers can be combined, provided that they make no errors and reason logically. This then implies that there* is *a kind of "global state," in the limited sense of a state including all available knowledge. If two observers' beliefs are so inconsistent that they cannot be combined together with any amount of communication and experimental data, they might as well be living in different worlds. This is why we say that insane people are "out of touch with reality."*

I can see why one would say that, especially if one believes that the process of science has an end, and that there is a sense in which the universe is pre-formed. (So maybe I will admit to *some* postmodernist tendencies.) But then ...

Brunism 5: *Also for this reason, a rational person may very well never assign a perfectly pure state to a system, but always a mixed state of the form*

$$\rho = (1-\epsilon)|\psi\rangle\langle\psi| + \epsilon\rho'$$

where ϵ contains the unspoken acknowledgement "But I might be wrong."

giving an agent the right to set his own density operator seems to me to be a concession to the quantum state's (purely?) epistemological content.

A note on the "The Second Amendment" to be sent to David soon after lunch.

15 August 2001, to Brun, "Re: Knowledge, Only Knowledge – Reprise"

Brunism 6: *Frankly, I find your negativity about this problem puzzling.*

On the contrary, I'm quite taken with it: otherwise, I wouldn't have given it the time of day. Reading the collective emails has been a great learning experience for me.

15 August 2001, to Schack, "Compatible States"

Schackcosm 2 : *I got a number of rather incoherent messages on compatible states from you. Is there anything more recent on this?*

Yes. About three thousand more notes. The issue refers to David Mermin's recent `quant-ph` paper "Whose Knowledge?"

Where it stands right now is with Brun, Finkelstein, and Mermin all ganged up against me (as being the unreasonable one). The conclusion they've come to is something like this:

> Two density operators ρ_a and ρ_b can describe the same system iff the support of ρ_a has a non-trivial intersection with the support of ρ_b.

I, on the other hand, think such a statement is far too dictatorial for a Bayesian's taste. Try as I might, they just don't think my points are relevant.

I'll paste (what I deem to be) my most lucid notes below. You can judge for yourself. I'm soon going to drop out of the debate, I believe: it's now at a point of diminishing returns. But, I would like to hear your thoughts!

16 August 2001, to Grangier, "Subject-Object"

Thanks for your note! I'm always amazed when anyone reads or skims my papers: You have a friend for life!

Last night I read your paper `quant-ph/0012122`, which I had never seen before. Thanks for bringing it to my attention.

Grangierisme 1 : *You are probably aware that statements such as : "The quantum state is information. Subjective, incomplete information." ... "Quantum states are states of knowledge, not states of nature. That statement is the cornerstone of this paper." are unwarranted, since just opposite statements can be made without changing any physical predictions or even any technical development. My personal view is that these statements are even wrong, as soon as "quantum state" is understood as "pure quantum state" (see eg* `quant-ph/0012122`*, missing in your list on p.1).*

In your words,

> [C]ontrary to the copenhagian dogma, a central point in our approach will be to give an "objective reality" to the quantum state of a physical system, in a sense which is developed below. ... The quantum state of a physical system is defined by the values of a set of physical quantities, which can be predicted with certainty and measured repeatedly without perturbing in any way the system.

Here is a problem I have with this conception. (It is a problem I am quite sure you are aware of, but for some reason you did not address it directly in your paper.) Consider two electrons originally prepared in a spin-singlet state—one electron in the possession of Alice, one in the possession of Bob. Let us imagine now two alternative scenarios. In one, Alice measures σ_x on her particle; in the other scenario, she measures σ_z. By your criterion, Bob's particle does not start out with a quantum state (since the two electrons are in an entangled state)—which is fair enough (I have no qualm with that)—but *immediately upon* the measurement it will go into a definite quantum state, either an eigenstate of σ_x or an eigenstate of σ_z, depending upon the scenario. We know this because if Bob were to thereafter repeatedly send his particle through a Stern-Gerlach device

with the proper orientation (for the given of the two scenarios), Alice would be able predict with complete certainty which way Bob's particle would go. Moreover, if Bob is careful enough, these further measurements on his part will not perturb his system (in the sense of changing the spin quantum state of the electron). So, we have just what you had wanted: complete certainty and no necessary perturbation.

But if the quantum state is an objective feature of the electron, then we see that it can be toggled one way or the other instantaneously from a distance. (Alice's measurements causes Bob's electron to go into one or another quantum state instantaneously.) Thus, if you accept the objectivity of the quantum state, then you must also accept the objectivity of instantaneous action at a distance (that in no way diminishes with distance or the particulars of the medium filling the space between Alice and Bob).

This is something I'm not willing to do. Not out of dogma, but because it strikes me that the world would have to be horribly contrived to have this property: a little private (instantaneous) telephone line between each and every physical system that will not accept outside calls. I.e., I can never make use of this instantaneous action for real, live communication even though it really, really, really is there. It stretches my imagination too far.

You, of course, may accept it as you wish; but the reasons above are mine for not doing so. Let me address your other point that I quoted above:

> [Your statements] are unwarranted, since just opposite statements can be made without changing any physical predictions or even any technical development.

I certainly disagree with the latter part of this sentence: that was the whole point of my paper (i.e., that it can change the technical development of the theory). Taking one or another point of view about the objectivity of the quantum state motivates different directions of theoretical exploration. In my case, it motivated trying to find the four theorems I presented in that paper. One who believed in the objectivity of the quantum state—I am quite sure—would not have sought those theorems in the first place.

Grangierisme 2: *By the way I also disagree with your point of view that "Quantum Theory Needs No 'Interpretation'," Phys. Today* **53***(3), 70 (2000). The fact that a physical theory ALWAYS needs an interpretation is in my opinion a central difference between physics and mathematics.*

You won't find a disagreement with me here. The title and closing sentence of that paper were meant to be tongue-in-cheek plays on something Rudolf Peierls once said: "The Copenhagen interpretation *is* quantum mechanics." The whole paper is very definitely about an interpretation, and why one does not need to go any further than it to make sense of quantum mechanics as it stands. My paper `quant-ph/0106166` and the large (more personal) collection `quant-ph/0105039` is about going the next step, i.e., what to do once we have established the belief that quantum states are states of knowledge.

When we do finally dig up an ontology underneath quantum mechanics, I'm quite sure it will be an interesting one!

17 August 2001, to Caves-Schack, "Unloading Bayes"

Let me unload a couple more Bayesian thoughts on you—i.e., some things we will probably want to address in the RMP article.

1) Attached below is a note I wrote to Mermin giving what I think is the cleanest justification for noncontextuality in any Gleason theorem. In fact, it shows that noncontextuality is more basic

in the hierarchy of theories than anything else we've dealt with yet. I.e., it comes far before the particular details of quantum mechanics. (Maybe this is what Carl has been saying all along, but I had to work through it for myself before it stuck.)

2) Let me bring your attention to a cluster of papers that I think are really important.

> quant-ph/0010033 [abs, src, ps, other] :
> Title: Quantum computing via measurements only
> Authors: Robert Raussendorf, Hans J. Briegel (LMU Munich)
> Comments: 4 pages, 2 figures
>
> quant-ph/0108020 [abs, src, ps, other] :
> Title: Universal quantum computation using only projective measurement, quantum memory, and preparation of the 0 state
> Authors: Michael A. Nielsen
> Comments: 4 pages
>
> quant-ph/0108067 [abs, src, ps, other] :
> Title: Computational model underlying the one-way quantum computer
> Authors: Robert Raussendorf, Hans Briegel
> Comments: 26 pages, 3 figures

I am especially taken with the Raussendorf-Briegel development. On one level, I think it might be the simplest avenue for addressing quantum computation from a Bayesian point of view: One just builds up the proper initial (universally valid) entangled state, and then does a measurement site by site, doing proper (Bayesian) updating of the quantum state of the remaining sites at each step. At the end of the day, one's knowledge is updated to the sought after answer. On another level (not for RMP), I think it starts to capture what I have been hoping for as an explanation of the power of quantum computation: it is not quantum parallelism that is doing the trick, but the "zing" of quantum systems that makes them sensitive to our interventions.

3) Let me put further below my replies to Khrennikov on his "contextual probabilities" business. I've already shared this with Carl, but not Ruediger yet. What I like about my statement there is that it starts to put Gleason's theorem in a more Dutch-book kind of light. Just as in classical theory, the setting of initial probabilities is completely free (and therefore subjective). What is set by coherence/rationality is the transformation rules. To that extent, this is the objective content of probability theory. Similarly with quantum mechanics in the light of Gleason's theorem: The objective content of quantum mechanics (or at least part of it) is that if we subjectively set our probabilities for the outcomes of any informationally complete POVM, we are no longer free to set them arbitrarily for any other observable. The probabilities are now fixed and can be calculated uniquely from our previous subjective judgment. (I hate the American way of spelling "judgement.")

That's about it for now. I've certainly written loads more (Bayesian oriented stuff) since we've last seen each other. But nothing else that might have been striking is coming to mind right now. I'll probably unload more as I think of it.

As I'm hearing about more and more of you arriving in Santa Barbara, I'm starting to get a little envious that I'm missing all the fun. At least Kiki and Emma will be off to Munich Monday evening. I won't be joining them until the following Wednesday. I have this great dream of idling away my hours in the mean time with Arthur Schopenhauer, William James, and the Reverend Bayes. But I'm sure some reality will set in to knock me off my course of purity.

20 August 2001, to Caves, "Re: Unloading Bayes"

Cavesism 2: *I'm not sure anyone is going to be convinced by the above, but I think it is a start. So here's the scenario. There is a big set of things that can be true or false. The big set is determined by the (physical) theory you are dealing with. There are further rules that say how the elements in this big set can be gathered into subsets that correspond to questions whose outcomes are exclusive and exhaustive. Now, if you make noncontextual probability assignments to the questions, you have just ignored the structure that led to your original set. That being the input from your basic theory, if you make noncontextual assignments, you are deciding not to pay attention to your own theory. Not too bright. So you should make noncontextual assignments.*

I know this won't come as a surprise to you, but (really) I found my (operational) explanation much more convincing. So, I guess I really didn't take my cue from you after all. Probably more from Pitowsky/Renes/Hardy.

Can you pinpoint what you didn't find convincing about my argument?

My trouble with your argument is that I still don't find the statement, "if you make noncontextual probability assignments to the questions, you have just ignored the structure that led to your original set," all that compelling. I guess I still don't quite understand what you are saying.

21 August 2001, to Grangier, "Contextual Reality = Information??"

Thanks for the longer explanation. It has indeed clarified things for me.

What I think is funny though, is that for precisely everything you say below (which I quote), I would call the quantum state "information" rather than an "objective reality."

Let me ask you this: Would you reread Section 3, "Why Information?" in my `quant-ph/0106166` and comment why and how you would use a different language in rewriting that? It seems to me that to say "a (pure) quantum state is objective, but it is contextual, ie, defined relatively to a particular set of measurements results" (as you do), can only be to teeter very close to admitting that the state is information (and only information). If the state is only defined relative to Alice's measurement results, why not call those measurement results the actual reality and be done with it? What does having the quantum state being objectively real add to the story?

Grangierisme 3: *Alice may obviously choose between several measurements, but she must eventually decide for one and perform it. My definition of the "objective quantum state of a system" REQUIRES that Alice's measurement is completed, which does not "toggle" the state, but simply define it (at Alice's location).*

Once Alice's measurement is completed, and given her measurement result, she will be able to make predictions about Bob's state. But obviously the corresponding information (orientation of Alice's polarizer and measurement result) will have to travel (non-instantaneously) from Alice to Bob. Thus there is nothing here like "objective instantaneous action at a distance", that I dislike as much as you do. On the other hand, there is an objective quantum state, in the sense that once Alice's results have reached Bob (and not before !), Alice and Bob will know "the values of a set of physical quantities, which can be predicted with certainty and measured repeatedly without perturbing in any way the system". In this view BI are violated as a consequence of the lack of "separate reality" of the two particles, rather than as a consequence of non-locality.

To put it another way : a (pure) quantum state is objective, but it is contextual, ie, defined relatively to a particular set of measurements results. If new measurements are done that are not

in the initial set, the definition of the new physical state MUST include the results of the new measurements. In some sense, this is simply a restatement of Bohr's 1935 answer to the EPR article (and maybe this can be related to your "Bayesian" approach ?).

21 August 2001, to Caves, "Noncontextuality"

Cavesism 3: *This is a weird notion of operational, since you rightly note immediately afterward that it is really tautological. How does it justify assigning the same probability to an element in the big set, no matter what the context, to note that you have a rule that assigns the same probability to an element in the big set, no matter what the context? It seems to me the real point is why you would ever identify two elements from different contexts, and the reason is that that they are the SAME element in the underlying set, which is handed to you by your underlying theory. Moreover, you only have one thing to go on, and that is the fact that the theory tells you that two elements are actually the same element in the underlying set. This means that the theory wants them to have the same probability in all contexts.*

I think what is creeping in here (anew) is a fundamental philosophical difference that is starting to come between us. Maybe I can characterize it in the following way. You want the theory to come first, and then to (somehow) recover our activity as scientific agents back out of that grander picture. I, on the other hand, am becoming more and more content to start with the scientific agents and thereafter pluck out those terms in their discourse with various common features to call a theory.

Cavesism 4: *How does it justify assigning the same probability to an element in the big set, no matter what the context, to note that you have a rule that assigns the same probability to an element in the big set, no matter what the context?*

It doesn't justify assigning the same probability. Assigning the same probability is the very reason for assigning a common element in our theoretical descriptions of two very distinct devices.

22 August 2001, to Grangier, "Contextual and Absolute Realities"

Grangierisme 4: *This may sound rhetoric, but if a theory explains nothing less than the stability of matter, is able to calculate g-2 (and many other things) with an incredible accuracy, and nevertheless does not speak about "physical reality", what is physics? I thus consider much more useful to put forward a reasonable definition of reality, that allows me to say to a journalist : "QM speaks about the reality of micro-objects. It is a weird reality, ... but it is REAL."*

You ought to know that I could not agree more with this (the part that I quoted, not the part that I did not quote). The issue on my mind is whether it is productive to view the quantum state in particular as THE term in the theory that corresponds to the objectively real. The way I view the issue is this: Quantum theory is a mixture of objective and subjective elements and we will only make progress in quantum foundations when we have had the intellectual strength to cleanly separate those two ingredients. With this point of view, I can answer the journalist just as soberly as you (i.e., without carrying him off to a postmodern fuzz fest).

There are certainly elements in quantum theory that I am immediately willing to identify as objectively real. A good example comes from our Alice and Bob example, generalized ever so

slightly. Suppose Alice's system has a Hilbert space of dimension E and Bob's has a Hilbert space of dimension D, and that again Alice and Bob start off with an entangled state for the bipartite system. I say that the quantum state must be (freely) subjective information because, depending upon what measurement Alice chooses to perform on her side, she will ascribe one or another quantum state to the system in Bob's possession. However, there is one obvious thing that Alice cannot change by the free choices she makes on her side of the world: It is the dimensionality of Bob's system's Hilbert space. Thus, I would say the number D is the objective reality in this situation. The number D remains constant regardless of what Alice does.

I could put this in language more to your liking, saying something like:

> The quantum state that arises for Bob's system after Alice's measurement is a "contextual reality," whereas the Hilbert-space dimension of Bob's system is an "absolute reality."

but I personally don't see that as a road to further progress. I.e., it distracts from what I view as the main point of our task in clearing up the foundations: namely, working hard to separate the subjective from the objective.

What I want to know in the most physical of terms is what does that number D signify? I want to find a way of describing its meaning that never once refers to a quantum state. When I can do that, then I will say that some progress has been made in identifying the objective part of quantum theory. But that is just an example.

Grangierisme 5 : *Remember my quotation about the 2 electrons in He : in their singlet state is "subjective information", how do the electrons know it?*

It seems to me, this is just a varied form of the Penrose argument I wrote at the end of my Section 3. I dismiss it in the same way that I did there. From my point of view, to say that the electrons are in a singlet state is to give nothing more than a compendium of things we can say about how they will react to our experimental interventions into their nonhuman bliss. It is not they who can predict the consequences of my invasions into their territory, it is me.

22 August 2001, to Caves-Schack, "Identity Crisis"

I enter this note with trepidation, because I know that what I am about to say will not be taken lightly by either of you, and chances are you will just view me as a troublemaker (again). I don't want to be a troublemaker, but I do have some concerns that are starting to eat more and more at me.

The problem is, I am starting to have serious misgivings about our Sections II and IV of "Making Good Sense" (PRA Version). Most of this new thinking has come about through my taking David Mermin's quest in his paper "Whose Knowledge?" to task, and my ensuing debate with Todd Brun, Jerry Finkelstein, and David himself. But some of it, I think, flows directly from the spirit of Bruno de Finetti, which I now believe I had shut my eyes to for too long.

The issue is no less than whether we really believe probabilities are subjective or not. I think a failure on our part to take their subjective character completely seriously is causing us to go down a path I'd rather not take.

Let me try to explain as best I can. The trouble is localized in our claim:

> ... if two scientists have maximal information about a quantum system, Dutch-book consistency forces them to assign the same pure state.

In the mildest version of my troubles, I am starting to think this statement is contentless. In the stronger versions, however, I find it misleading, and I almost want to say "wrong" (though maybe I won't go that far).

To make sense of what I mean by this, let me start by taking a cue from Bernardo and Smith:

> Bayesian Statistics offers a rationalist theory of personalistic beliefs in contexts of uncertainty, with the central aim of characterising how an individual should act in order to avoid certain kinds of undesirable behavioural inconsistencies. ... The goal, in effect, is to establish rules and procedures for individuals concerned with disciplined uncertainty accounting. The theory is not descriptive, in the sense of claiming to model actual behaviour. Rather, it is prescriptive, in the sense of saying "if you wish to avoid the possibility of these undesirable consequences you must act in the following way. — p. 4, Bernardo and Smith, *Bayesian Theory*

If we accept this, then I think there is a much better way to word our "$p = 1$ when certainty" addition to the Dutch book argument in Section II. It seems to me it should more properly be viewed as a normalization condition, subordinate *only* to internal consistency/rationality as all the rest of the Dutch-book argument is. I say the latter to contrast it with how we presently have the argument worded in our paper—namely, by making the $p = 1$ condition subordinate to some objective feature of the world. We write:

> The only case in which consistency alone leads to a particular numerical probability is the case of certainty, or *maximal information*. If the outcome E is certain to occur, the probability assignment $p < 1$ means the bettor is willing to take the side of the bookie in a bet on E, receiving an amount px up front and paying out x if E occurs, leading to a certain loss of $x(1-p) > 0$. Consistency thus requires that the bettor assign probability $p = 1$. More generally, consistency requires a particular probability assignment only in the case of maximal information, which classically always means $p = 1$ or 0.

What does it mean for an outcome E to be "certain to occur"? I think that phrase is much more loaded than we have previously treated it. In the Dutch book argument there are three players, two of them animate and conscious (the bettor and the bookie) and one of them presumably inanimate and unconscious (the world). To which player does the certainty get attached? I don't think we make this clear in the way we ought to.

If the certainty is to be attached to the world, then what business does it have to do with my subjective judgments (which by definition cannot be in a bijective correspondence with the world's states)? Instead, I would say the "certainty" can only be a subjective judgment in and of itself. The Dutch-book argument for requiring $p = 1$ in the case of certainty should then be more accurately advertised as a call to be "true to our hearts." I.e., the argument is really that, *when* we believe an event will happen with certainty (a nonnumerical judgment), then we should ascribe $p = 1$ (a numerical judgment) for booking purposes. That is, Dutch-book coherence gives us a way to translate a nonnumeric belief into a numeric one.

The thing that is really at issue here is that I think we should remind ourselves always that "certainty" itself is nothing more than a belief. It may be a belief that can ultimately be tested against the world in a single shot, but nonetheless it is a belief. I believe with all my heart that my mother loves me; Schopenhauer believes with all his heart that she hates me. The only thing Dutch-book consistency can give us is that "if we wish to avoid the possibility of undesirable consequences" (Bernardo/Smith), then I should ascribe $p = 1$ and Schopenhauer should ascribe $p = 0$ to the proposition "love." The Dutch-book argument prescribes that we each should be true

to our hearts—that we should both act in accordance with our beliefs. But it does not have within it the power to make us believe the same thing ... EVEN in the case of "certainty."

Now, we whitewash all that by introducing the phrase "maximal information," which somehow makes "certainty" seem more objective, but now I'm starting to think that that phrase is pretty impotent in this context. What role does it really play in our argument? I can't find any, other than that it is a euphemism for declaring that we *believe* we have nothing left to learn (in the sense that we *believe* there is nothing left to learn from the remainder of the world that will help us refine our predictions for the system at hand). That belief may be wrong in the sense that rationality AND the world will not allow us to perpetuate the belief AFTER the experimental trial, but until the trial, "certainty"—from the Bayesian view—can mean nothing more than a metered belief.

You should be able to tell where I'm going with this by now. In Section IV, we write:

> Maximal information in quantum theory instead corresponds to knowing the answer to questions that share one particular projector.

I suppose what I am saying is that I just cannot accept this anymore. At least not in its present form. Instead, if I were to modify it to bring it into alignment with everything I said above, I would have to write something like this:

> Maximal information in quantum theory instead corresponds to believing adamantly that one knows the answer to all the questions that share one particular projector.

You might think this is nitpicking, but it completely takes all the steam out of Section IV. For it gives Dutch-book consistency no grounds for enforcing that two agents "with certainty" should believe the same thing. And consequently it gives no grounds for enforcing that two agents with "maximal information" should make the same quantum-state assignment.

The only way I see to reinstate our original role for maximal information is to say that two observers can only have maximal information *when* they are both right (in the sense that the world MUST CONFORM to their probability one predictions). But then, using our argument for Gleason's theorem, we would have DERIVED that quantum probabilities are objective probabilities! (This will be my only exclamation point in the whole note, so you should take it in seriousness.) That is, we would be saying that we have maximal information only when we *know* an objective reality, and by our derivation, that objective reality would then be equivalent to a compendium of probabilities.

Instead I think the best we can say is: If Alice and Bob both believe adamantly that they know the answer to some potential measurement AND that measurement happens to be the same for both of them, then Dutch-book consistency and Gleason's theorem will enforce that they make the same probability assignments for all other measurements (i.e., that they assign the same quantum state). But said that way, I don't think any non-Bayesian will be particularly impressed: For they would say that all we have shown is "if two people know the same things, then they will know the same things." Woop-ti-do.

The Bayesians among us will still have some room to be impressed: For it will not be a priori obvious to them that beliefs about one observable should have anything to do with beliefs about another. In particular, it might even surprise them that a common belief in certainty (for two observers) for any fixed observable should lead to equal probability assignments for all other observables. But even then, I think our the shock-value of our paper will be diminished. For I think in no way have we shown that when two observers make two pure-state assignments for a system, those pure states MUST be identical.

For me, this is a liberating thing to understand, i.e., that there are no dictatorial constraints on quantum state assignments. But I suspect you will feel otherwise, at least on the first reading of this note. So, let me beg your forgiveness in advance.

As I alluded to in the beginning, these thoughts of mine don't live in a vacuum. They have been spurred by my debates with Brun, Finkelstein, and Mermin. Thus let me give you some more material to chew upon: I'll attach it below, in the form of a composite note that I've already sent to Ruediger. Perhaps it will help clarify the things that have brought me to this position. Some of it, of course, will require that you try to imagine the context, but I think the notes are self-contained enough that you will be able to fall into the line of thought and see its relevance.

Best wishes (in spite of my predictable trouble),

23 August 2001, to Caves, "My Own Version of a Short Note"

I'm just back from a very long day in NY City (bookshopping), and a very long night before that (reading). So I won't reply to your notes until I'm a little more refreshed. But one quick comment:

Cavesism 5: *I get the feeling, strengthened by your own confession that it is true, that my e-mail doesn't make much of a dent, so why bother with it.*

That's absolutely not true. I read everything you send me many, many times over. When they are reasoned well, I accept your arguments. And you know I much prefer this method of communication, just so I can have the opportunity to fully understand what my correspondent hopes me to absorb—I've never been a quick thinker, and this helps me fill in for that inadequacy.

Today I focused on rounding up some more William James, John Dewey, Percy Bridgman material. I think James is taking me over like a new lover. I had read a little bit of him before, but I think I was more impressed with his writing style than anything. But I was drawn back to him by accident, after reading Martin Gardner's *Whys of a Philosophical Scrivener*. Gardner devoted a lot of time knocking down James' theory of truth, because it is just so much easier to accept an underlying reality that signifies whether a proposition is true or false, rather than saying that the knowing agent is involved in eliciting the very proposition itself (along with its truth value). And something clicked! I could see that what James was talking about might as well have been a debate about quantum mechanics. He was saying everything in just the right way. (Let me translate that: he was saying things in a way similar to the way I did in my NATO "appassionata.") And things have only gotten better since.

Have safe trips to everywhere you need to go. Kiki is due December 23, but I'll see what can be done about ITP in November.

27 August 2001, to Mermin, "Bayesian Pill Taking"

I am finally writing to reply to your long note on "knowledge." I apologize for keeping you waiting so long, be even now—since Ruediger is still in the woods—I feel a little like I am writing you prematurely. Nonetheless, I have this overpowering desire to get this issue out of my mailbox and be done with it. So here, I am. I would ask you, however, to please keep these thoughts private, at least until Ruediger finally emerges. For, what I am about to say involves him (and Caves) directly.

I know that you think your note is ultimately conciliatory, writing:

Merminition 8: *It seems to me that these are all valid statements about the formal structure of quantum mechanics, independent of what interpretation you favor.*

But after much soul-searching, I still cannot agree with your language. The soul-searching was required because this position of mine flies in the face of some of my very own published words (`quant-ph/0106133`).

The difficulty hinges on your Proposition 1:

Merminition 9: *1. A system that is known (by somebody) to be in a state psi cannot be found (by anybody) to be in a state orthogonal to psi.*

In my new view, to say this is to throw away all that we have been striving so hard for in establishing that quantum states are subjective entities (and purely subjective entities). You admit so much yourself at the end of your Comment (c). [You can take this to mean that I also disagree with the beginning part of Comment (c).]

In 1880, I suspect there was not a single educated physicist who doubted one iota that the speed of light was set with respect to the stationary ether. Regardless of that, in 1881 the first evidence came out against the common belief. The *spirit* of your Proposition 1 would outlaw the very happening of that wonderful historical event.

I see no way (nor even want to anymore) to get around this. If the quantum state is a subjective entity, without rigid connection to the world in itself,—as I think Einstein's original argument for its subjectivity indisputably shows us—then this is something we simply have to live with. We must allow that experimental data can speak against our predictions, no matter how set we are to believe that they will not. Yet, on top of this, we must also accept that there is no objective sense in which a (pure) quantum state is simply "wrong" before the experiment is performed.

I put all my heart and soul into presenting this point of view in the note below to my coauthors, and I will share it with you. It took me a whole day to write the thing; I would like to think it is worth reading. (Also I have slight evidence that it makes sense in that Caves ultimately concurred himself.) I believe the note is just as relevant to you and your note on "knowledge" as it is to Caves and Schack. I will also send you the PRA version of `quant-ph/0106133`, so you can track the sections.

There. Now I will clean out my mailbox.

27 August 2001, to Caves, "A Little Contextuality on Noncontextuality"

I've finally got enough time to write a small reply to your old note.

Cavesism 6: *There are two distinct approaches here, and I don't know which is to be preferred.*

1. The first point of view, which I have been pushing (as a way to justify noncontextuality), is that there is an underlying theory that sets up the structure of questions. This theory is primary. Noncontextuality emerges as the natural assumption that probability assignments should recognize the structure provided by the underlying theory.

2. The second point of view, which you have been pushing, starts with the role of scientific agents, as you put it, and uses the fact that probabilities are the same in different contexts to say that the elements with the same probabilities in all contexts explains why they are actually the same element in different contexts.

I do want you to understand my position, which is that I appreciate both these points of view. I'm not sure which will be the most fruitful in the long run. But they are trying to do quite different things. Here are two points:

1. The second point of view doesn't provide a justification for assuming noncontextuality, as you understand. Coming at things from the back door, as it were, it uses the fact of noncontextual quantum probability assignments to conclude that apparently different things are, in fact, the same.

I would not use language quite like that. I would say it IS a justification for noncontextuality. And it relies on quantum mechanics not one iota. Noncontextuality should be a property of any instrumentalistic theory (where Bayesian probability has been grafted onto to the world as the best way for us to steer our actions within it). By an instrumentalistic theory, I mean one where we explicitly have to talk about our various possibilities for experimental intervention into nature—a theory where we cannot detach the experiment from the phenomenon.

Cavesism 7: *2. The first point of view appeals to me presently because it manages to make a long straight run to the state-space structure and the quantum probability rule given only the Hilbert-space structure of questions and probabilities faithful to that structure (i.e., noncontextuality). I think the second point of view needs to address the following question: given a set of elements to which noncontextual probabilities are assigned, what structure is forced onto the set by the existence of these noncontextual probabilities? This question seems hopelessly underconstrained to me, but Howard described to Joe and me some math research on this sort of question.*

What you say "needs" to be done, seems hopelessly underconstrained to me too. But I think you shouldn't view the problem that way. The point is, one simply has a way of clearing the air of the noncontextuality issue BEFORE getting down to the nitty-gritty of quantum mechanics. Noncontextuality is the base assumption upon which one plays a new game: What *physical* assumption makes it so that our instruments should correspond to POVMs and not some other mathematical structure?

By the way, you know I really dislike your phrase "given only the Hilbert-space structure of questions." I've probably said this to you before, but let me try to articulate why in more detail so that maybe you'll remember it a little better. In your own words, the phrase is "hopelessly unconstrained." What does it mean? It seems to me there are all kinds of possibilities one could imagine if one didn't know quantum mechanics beforehand. Here's a simple example: An even more basic feature of Hilbert-space before orthogonality is linear independence. When God came down and said, "You will use Hilbert-space structure for the questions you can ask of nature," why did he not mean that any set of linearly-independent vectors corresponds to a valid question? Presumably there are good reasons. But those good reasons need to be spelled out, and are not at all implied by the simple phrase "Hilbert-space structure."

Let me send you to pages 86–88 and pages 361–362 of the samizdat. There it is shown that linear independence does not mesh so well noncontextuality. It is a dumb exercise, I know, but it does indicate that ones need to be careful with one's phraseology.

Oh, let me tell you another thing, of historical note. I talked to Howard Friday, and he tells me that this kind of justification for noncontextuality goes all the way back to Mackey (but then he settles on the setting of ODOPs thereafter).

Cavesism 8: *I'm also quite interested in another question: How general are the theories where maximal information does not lead [to] certainty, yet does lead to unique noncontextual probability assignments? In other words, for what classes of theories is there a Gleason theorem?*

I think this is a really good question. I was talking to Eric Rains the other day and he thinks that the appropriate generalized setting might be the Jordan algebras. This is because this is the largest structure he knows where there is a notion of positive operator. (I had shown him the trivial POVM way of proving Gleason.)

If Gleason can be proved in such a wildly general setting, I think it would be quite interesting. For it would tell us that the quantum probability rule is not very closely connected to physics at all. Dreaming of the process of deriving quantum mechanics as successively tucking up the more general structure of Bayesian probability theory, one might say that the real physical assumptions don't come until much later in the game. That would be worthy knowledge!

Waiting for a stupid doctor's appointment,

28 August 2001, to Mermin, "Re: some questions re your comments"

Merminition 10: *The trouble with our exchanges is that I'm always trying to zoom in on the issues under dispute and you're always trying to zoom out.*

Thanks for the note. You're point is a good one. (It's just that I've had to write so much email lately, and always, that it does take a toll. For instance, the present debate has been particularly taxing in that regard. Yet the issues have been important enough to not give up. So certainly I was hoping to recycle some material.)

I will try to write you a (focused) reply tomorrow, after recovering from today. Tomorrow evening, I fly out to Munich for a week and a half. Then I go to the quantum foundations meeting in Ireland to tote the wares.

Still tonight I've got to work on packing, etc.

01 September 2001, to Mermin, "Left Wing, Right Wing, Not a Wing to Fly With"

I was going to use my day today to write you a long, thoughtful (but focused) note on all the recent issues you raised with me on "Whose Knowledge?", but now you've gone and angered me. I mean that.

Merminition 11: *I took my left-right terminology from the Science Wars. It seems to me that in arguing against anything objective other than knowledge Chris is taking a decidedly post-modernist position and therefore allying himself with the "Academic Left" attacked by Gross and Levitt.*

Is my point of view so very subtle that even my most sympathetic patrons cannot decipher it? Or have I finally caught on that you're really not listening to me after all?

You can't stand this, but you deserve it: I will excerpt part of a note I recently wrote to Andrew Landahl. What set me off in his case was when he wrote me the following after having read my paper "Quantum Foundations in the Light of Quantum Information" and given a Caltech journal club talk on it.

> On the whole I portrayed your "party platform" as the statement that "Quantum states are states of knowledge about the consequences of future interventions." In particular, those consequences aren't consequences to reality, but rather consequences to

states of knowledge about even further future interventions. In this worldview Bayesian agents don't work to align their predictions with an underlying reality. Instead they work to align their predictions with *each other*. It is as if reality in this picture is solely the agreement of predictions!

I'd be interested to hear if you believe that this is a fair characterization of your party's platform. After reading this paper, I came to the conclusion that you didn't believe in reality at all. (Or at best I thought you believed reality = knowledge.) John Preskill tells me you believe otherwise, namely that there *is* a reality, which surprised me. I'm curious to hear what you believe reality is.

The scheme below is that everything marked with a ">" is a direct quote from Andrew Landahl's letter. Everything else is either me, or a quote from my paper. I will put only the very most relevant part of my reply.

Please do read it before you—YES, you, the most trusted of my academic friends—slander me again. It is nothing if not EXACTLY relevant to what you wrote about me above.

With surprisingly kind regards,

02 September 2001, to Mermin, "Re: Left and Right"

Merminition 12: *I thought you were safely away in Ireland or pre-Ireland.*

Yes, I am in pre-Ireland mode (in Munich), but keep in mind that I am never safely away.

Merminition 13: *I'll read what you sent in a little while. But note that I would not have made such a remark (even in jest) before I got your comments on my summary (and the cc of your letter to your own coauthors) which struck me (and I thought you too) as going beyond your earlier position.*

I will take this remark into account for the longer reply I am presently constructing for your earlier query. (It is sitting at 6K in length now, and will likely be finished tomorrow. Right now, I'm having my first beer of the evening.) But preliminarily, let me say that the only thing that letter to my coauthors did was clarify my position on the subjective character of the state vector. In that I went further than before, having gotten weak in the knees briefly about my position on two agents sharing differing quantum states. But, I do believe that I have long held fast in my opinion that there is something in the universe beyond human ken: It has always been a problem of finding the right language and right ideas for expressing what that something is, AND how it is PARTIALLY reflected in the structure of quantum mechanics.

If this does not make sense yet, I hope it will make more sense after my long note to you tomorrow.

02 September 2001, to Mermin, "Intersubjective Agreement"

Let me finally throw myself into the ring of intersubjective-agreement to see if I can wrestle you down a little. I will try to be every bit as focused for you as the issue will let me be.

Merminition 14: *Well maybe you're more radical than I thought. It was to avoid correlations floating in the void, unattached to anything whatever, that I've been interested in trying to follow you down the path of knowledge. If all it led to were knowledge, floating in the void, unattached*

to anything whatever — even to other knowledge — then I'd be no better off [than] when I started down the trail.

I think you misunderstand something very deeply here. The point of separating the categories "knowledge" and "reality" (or "subject" and "object" for that matter) is not to make knowledge an objective reality in its own right or, even worse, to make it the sole reality. Rather it is to say that there is a distinction and that that distinction should be recognized. I see nothing wrong with allowing a physical theory (such as quantum mechanics) to contain formal elements that correspond to *both* categories. The issue in my mind is which elements should be thrown into which category? The answer is not completely clear to me, but I am fairly convinced of one thing: The state vector should not be thrown into the reality side of the line.

What I have ultimately NOT been able to stomach about your wording of the whose-knowledge "answer", and Jerry's wording of the whose-knowledge "answer"—some of Todd's versions would actually survive—is that you say, under certain circumstances, two scientific agents (observers, or what have you) MUST assign "consistent" quantum states to a given system. In the case of pure states, the two agents MUST assign the *same* pure state to the system.

Let me get that through your head: What I object to is the word MUST. Todd once wrote it this way,

Brunism 7 : *We have been describing a consistency criterion.* If *one wishes to combine two state descriptions of a single system into a* single *state description, the criterion tells one* when *it is consistent to do so (i.e., when the two descriptions are not actually contradictory).*

I agree that nobody is holding a gun to Alice's head and forcing her to incorporate Bob's information.

and to this way of speaking I can agree. But if you take away Todd's "*If*", then everything collapses in my mind. Enforcing that two agents MUST make the same state assignment if they are going to be "right" at all reinstates the very objectivity, the very agent-independence of the quantum state that the Mechanica-Quantica-Lex-Cogitationis-Est program has been working so hard to exorcise. [As you know, we made a serious misstep in our `quant-ph/0106133`, but that will be rectified in the next edition.]

It is much like the old debate. Is materialism right? Or is it Berkeley's idealism that is right? Who cares, I say. Both philosophies are just simple samples of realism: They only disagree on the precise concept which ought to be taken as real, mundane matter or sublime consciousness. The way you characterize it above, one would think that the only fruit of the Mechanica Quantica program would be the RENAMING of a material reality into an ideal one—a shift more of emphasis, rather than anything of grit.

Merminition 15 : *Are you also unable to agree with the statement that a photon that is known (by somebody) to have just passed through a horizontal polaroid cannot immediately thereafter be found (by anybody) to pass through a vertical polaroid.*

I'm asking you about this concrete example of the general proposition because I can't tell whether you're objecting to the language in which I generalized it or whether you object to the statement about photon polarizations too. If it's only the former I'm happy to use less provocative phrasing. All I meant by "be in a state psi" was "has been found to be" in the sense I specified prior to making the objectionable statement. But I worry that you object to both statements. In that case you are walking a dangerous path, denying that one of the most elementary applications of quantum mechanics has a legitimate meaning.

Here is what you are losing sight of. In the Bayesian world, two agents must agree a little before they can agree a lot. Agreeing a lot is the currency they are seeking, but agreeing a little to begin with is not the limitation of their existence. I'll come back to this from a more positive angle in a minute, but let me tackle your particular question before that.

What does it mean for "a photon that is known (by somebody) to have just passed through a horizontal polaroid?" Presumably it means that a particular quantum mechanical test, a POVM, $\{E_b\}$ has been performed and one of the outcomes of that test has obtained—in this case, the label b is "photon passed through horizontal polaroid." Now, you ask, "immediately thereafter [can it] be found (by anybody) to pass through a vertical polaroid?" Implicit in that, you are thinking that the state transformation, subject to the measurement outcome, is

$$\rho \longrightarrow E_b^{1/2} \rho E_b^{1/2}$$

up to normalization.

Suppose you are the somebody spoken of above; and let me be part of the anybody. Now, let us say that I stubbornly insist that the state transformation is

$$\rho \longrightarrow U_b E_b^{1/2} \rho E_b^{1/2} U_b^\dagger,$$

where the U_b are some unitary operators, and in particular, when

$$b = \text{photon passed through horizontal polaroid}$$

we have

$$U_b = \text{horizontal} \rightarrow \text{vertical}.$$

There is nothing in quantum mechanics (as a theory) that can keep me from believing that, so long as the ONLY thing specified is the "measurement" $\{E_b\}$. The point is, let us suppose we disagree on how our beliefs should get updated upon the incorporation of a measurement outcome into our knowledge bank. [As an aside, notice the distinction here: b is given the lofty title of knowledge, whereas ρ is subjected to being a belief. I allowed myself to do that because I am assuming we agree on the $\{E_b\}$, even if not the state-change rule. You might say we need at least this much to get the problem off the ground.]

So, it boils down to this in more common language,

Merminition 16 : *Are you also unable to agree with the statement that a photon that is known (by somebody) to have just passed through a horizontal polaroid cannot immediately thereafter be found (by anybody) to pass through a vertical polaroid?*

And I say no, I cannot agree: I saw Hideo Mabuchi in his lab yesterday, and I saw that he inserted a really fancy polarizer into his lab bench, one with an intriguing optical coating that allows horizontal photons in, but has them come out as vertical ones. I insist that I saw him do that: There's not a doubt in my mind. You insist that he is an honest upright boy, and he would never do such a thing to confuse us. With equal tenacity, there is not a doubt in your mind. We disagree, and in the strongest of ways.

How do we put our disagreement to test in the context of photo-detector clicks? WE insert a vertical polaroid—this one, I assume, we both do agree on—behind the "horizontal" one and see what happens. Aha! You were right after all! Mabuchi really did use an honest-to-god von Neumann polaroid; the input photon never made it to the final detector. It wasn't the fancy-coated polaroid after all, but I could swear I saw him put it in.

By now, I know that you are thinking I have gone through a ridiculously long-winded and pedantic way of describing a triviality: that one of us made a FALSE assumption. Implicit in your question was the reasonable starting point—indeed, the one we use in all common discourse—that all the agents involved start from a TRUE state of affairs.

But what can TRUE and FALSE mean in a world where our only handle for getting at things are SUBJECTIVE quantum states? We get at the world through our beliefs and belief updates—that's the fundamental tenet for me. And in that light, the only thing a FALSE belief can mean—as I put it to Caves and Schack in the infamous email—is that rationality (i.e., the Bayesian laws of thought) PLUS the world (i.e., the detector clicks, whose meanings in the end are also set by subjective considerations) will not allow the believer to perpetuate any remnant of his initial belief after the experimental trial. Let me say that sentence again for extra emphasis:

> The only thing a FALSE belief can mean is that rationality (i.e., the Bayesian laws of thought) PLUS the world (i.e., the detector clicks, whose meanings in the end are also set by subjective considerations) will not allow the believer to perpetuate any remnant of his initial belief after the experimental trial.

But if that is the case, what is so overpoweringly evil about having a "false" belief? Why must two valid quantum scientists necessarily be aligned in their beliefs, even in the case of "true" and "false"? I remain hardened: I see no compelling reason for asserting that necessity. Indeed, such an assertion is antithetical to the idea that a quantum state is a compendium of subjective degrees of belief.

If you think TRUE and FALSE mean something more substantial than I just described, then you tell me what role they play in my life other than as a kind of shorthand for some characteristics of my belief updates.

In pointing out these deficiencies, I am not "denying that one of the most elementary applications of quantum mechanics has a legitimate meaning." I am coming nowhere near that. I am merely asserting each scientific agent's constitutional right to believe what he will—i.e., to carry about whatever quantum state he wishes—SO LONG AS those beliefs do not contradict the constitution itself.

It is the latter that, from my view, is the most essential point you have been missing. So, let me get to that directly. With this I can finally start to define a positive program.

Merminition 17: *I don't see that limiting "objectivity" to mean "complete and necessary intersubjective agreement" is abandoning your quest. Indeed, your Bayesian authorities say as much,*

> We can find no real role for the idea of objectivity except, perhaps, as a possibly convenient, but potentially misleading, "shorthand" for intersubjective communality of beliefs.

It must be the "necessary" that raises your hackles [. . .]

If anyone cannot see by now that it is almost solely the word "necessary" that raises my hackles, then they are not listening. You wrote at the beginning of your note that:

Merminition 18: *If all it led to were knowledge, floating in the void, unattached to anything whatever — even to other knowledge — then I'd be no better off [than] when I started down the trail.*

But that is just not true. We have gained a serious amount of positive knowledge from this exercise. It has allowed us to see much more clearly what is firm and what is squishy in quantum mechanics. The state assignments I would say are always squishy; the rules for updating them are not. To the extent that these rules fulfill an edict in the spirit of Bernardo and Smith,

> Bayesian Statistics offers a rationalist theory of personalistic beliefs in contexts of uncertainty, with the central aim of characterising how an individual should act in order to avoid certain kinds of undesirable behavioural inconsistencies. ... The goal, in effect, is to establish rules and procedures for individuals concerned with disciplined uncertainty accounting. The theory is not descriptive, in the sense of claiming to model actual behaviour. Rather, it is prescriptive, in the sense of saying "if you wish to avoid the possibility of these undesirable consequences you must act in the following way.

I would say we have identified an objective piece of quantum mechanics.

It is not that "physics as intersubjective agreement" requires that agents always agree, or, at least, that they must agree in certain limiting circumstances. Instead it is that there is a procedure in any given case for deciding whether two agents will move closer to agreement than not after looking at the world. Quantum mechanics gives us such a framework. It might have been otherwise: One can imagine a world so chaotic that any percipient beings which happen to arise in it would forever be in their own individual dreamlike states, never realizing that it is even possible to come to agreement with their fellow quixoticoids.

Maybe a good (but limited) analogy is this. Think of an electric potential function from which, by taking a gradient, we can derive the electric field. The potential itself cannot be the real stuff because of its ridiculous freedom that can be set freely from observer to observer. Instead it is the way the potential changes spatially that is what is of interest. That spatial change which is the common denominator of all the disparate potential assignments is the real, real stuff. Now think of the quantum state in the role of the potential, and the quantum structures of POVMs, the Gleason theorem, and the state-change rule in the role of the potential's gradient.

As long as you and I play according to the quantum rules for updating our beliefs in your example—you with your

$$\rho \longrightarrow E_b^{1/2} \rho E_b^{1/2},$$

me with my

$$\rho \longrightarrow U_b E_b^{1/2} \rho E_b^{1/2} U_b^\dagger$$

—who is to fault one of us for being irrational? We just have different beliefs about how a state ought to get updated in the particular situation. Neither one of us is taking the constitution to task; neither one of us are using a state-updating method that does not fall into the quantum mold.

The analogy of this with classical probability theory is that we both might agree on the probability for some hypothesis $p(h)$, but disagree on the joint distribution $p(h, d)$ for hypothesis and data. Learning the data and using Bayes rule, we will generally then come to two distinct posterior assignments $p(h|d)$. That nevertheless gives us no warrant for backtracking our opinion that $p(h, d)$ is just a subjective belief (as are all probability assignments). Instead it helps us see that the objectivity working in the background is Bayes' rule; it is our common denominator.

Quantum states—or through Gleason's theorem, nothing more than compendia of quantum probabilities—do not float in a void. They are tied together more tightly than any other probabilities hitherto ever found. I cannot assign probabilities for σ_x outcomes, σ_y outcomes, and σ_z outcomes at the same time as *independently* assigning them to the outcomes of any more exotic POVMs. In changing my probabilities for the outcomes of some potential new measurements (just

after a previous measurement), I had better tie all those probabilities together along the lines of the general form of the quantum state-change rule.

In this, we see something like a much greater deepening of the Dutch book argument. In the classical case, we find that we will bring havoc upon ourselves if we allow ourselves to freely assign $P(A)$, $P(B)$, $P(A \wedge B)$, and $P(A \vee B)$ all independently. All compendia of quantum probability assignments must be tied to the particular structure of quantum states and the quantum state change rules. You should be thinking of the firmament rather than the void.

I think that's all the really general remarks I had wanted to say. Let me now *briefly* address the remaining specific points in your notes.

Merminition 19 : *If you do indeed object to both, then the only reason I can see for it (and I agree that this does raises non-trivial issues) is that probability 1 and probability 0 statements are idealizations — that nothing in the actual world we inhabit can be said to be certain or impossible. In that case, of course, the support of any acceptable density matrix is the whole Hilbert space and there is no content left to the criterion. But since the theory does allow you to talk about certain or impossible measurement outcomes, I'm reluctant to declare that its illegitimate to consider them in trying to develop a better understanding of the theory.*

I hope that you can see by now that "probability 1 and 0 statements being idealizations" (i.e., states of belief that we none of us, even Job, are really ever in possession of) has nothing to do with my considerations. A belief is a belief. Rationality itself cannot infringe on what numerical value a belief ought to be. It is therefore perfectly legitimate to think about these idealized situations.

Merminition 20 :

> CAF Said: on top of this, we must also accept that there is no objective sense in which a (pure) quantum state is simply "wrong" before the experiment is performed.

[Do] you also require me to accept there is no objective sense in which a pure quantum state is simply correct, after the experiment is performed?

Yes. The ghost of Bruno de Finetti haunts us:

QUANTUM STATES DO NOT EXIST

And I understand that oh so much better now than I did two months ago.

Merminition 21 : *P.S. I also asked for clarification of your views on objectivity as nothing more than intersubjective agreement, which on the one hand you seemed to reject in accusing me of going objective in comment (c) but on the other hand you seemed to endorse in quoting approvingly your Bayesian gurus.*

Bernardo and Smith would have never held fast to a "necessity clause" like you seem to be doing. That puts a gulf of distance between your two separate uses of the phrase "intersubjective agreement."

Merminition 22 : *P.P.S. Just got a blast from the Eastern Front (Mohrhoff — cc'd to you, I believe). I have the funny feeling that you two, who are so far apart in opposite directions (knowledge-without-facts vs facts-without-knowledge), may yet turn out to be strangely similar in some respects.*

To the extent that I understand him, I myself don't believe this is likely. The direction I see for physics, and quantum mechanics in particular, was perhaps no better put than by William James:

> Metaphysics has usually followed a very primitive kind of quest. You know how men have always hankered after unlawful magic, and you know what a great part in magic *words* have always played. If you have his name, or the formula of incantation that binds him, you can control the spirit, genie, afrite, or whatever the power may be. Solomon knew the names of all the spirits, and having their names, he held them subject to his will. So the universe has always appeared to the natural mind as a kind of enigma, of which the key must be sought in the shape of some illuminating or power-bringing word or name. That word names the universe's *principle*, and to possess it is after a fashion to possess the universe itself. 'God,' 'Matter,' 'Reason,' 'the Absolute,' 'Energy,' are so many solving names. You can rest when you have them. You are at the end of your metaphysical quest.
>
> But if you follow the pragmatic method, you cannot look on any such word as closing your quest. You must bring out of each word its practical cash-value, set it at work within the stream of your experience. It appears less as a solution, then, than as a program for more work, and more particularly as an indication of the ways in which existing realities may be *changed*.
>
> *Theories thus become instruments, not answers to enigmas, in which we can rest.* We don't lie back upon them, we move forward, and, on occasion, make nature over again by their aid.

Mohrhoff, from what I can tell, sees a "block universe" (to use another piece of Jamesian terminology). It is a completed thought in the cosmic consciousness.

Good wishes, and I hope this document answers more questions for you than it raises. Now I've got to run to the biergarten again for a little oompah-pah.

02 September 2001, to Mermin, "Truth and Beauty"

Here's another passage from William James's *Pragmatism* that may help reveal a little more of my mindset.

> The truth of an idea is not a stagnant property inherent in it. Truth *happens* to an idea. It *becomes* true, is *made* true by events. Its verity *is* in fact an event, a process: the process namely of its verifying itself, its veri-*fication*. Its validity is the process of its valid-*ation*.

03 September 2001, to Nielsen, "Subject/Object"

Nielsenism 1 : *You may also be interested to hear that I'm engaged to be married :-)*

Excellent! This is only a joke partially, but lately I've been so taken with the idea that unions can give rise to things greater than those contained in the parts—thinking of quantum measurement, in particular, from this angle—I've thought about calling my view on QM "the sexual interpretation of quantum mechanics."

Many congratulations!

04 September 2001, to Caves-Schack, "Note on Terminology"

Thinking about it more, I would like to emphasize a point that was buried away as an "aside" in my recent note to Mermin titled "Intersubjective Agreement."

I am becoming more and more dissatisfied with the slogan "A quantum state is a state of knowledge, not a state of nature." The reason for this is that people tend to view the word "knowledge" as something that can be right or wrong, depending upon whether it is in direct correspondence or not with something in the external world. For this reason—as brought out clearly in my debate with Mermin, Brun, and Finkelstein—I think we should get more into the habit of calling a quantum state a state of BELIEF. This is more in line with the language both de Finetti and Bernardo and Smith use for probabilities anyway, and therefore gets us into a quicker connection with the personalistic Bayesians.

I now think it is much better to reserve the word KNOWLEDGE solely for the outcomes of quantum measurements once they become part of the mental makeup of an agent interested in them. I walk into Mabuchi's lab, and to the extent that he and I agree that he is performing some POVM (denoted by a set of positive operators $\{E_b\}$), it seems to me valid to call the outcome b we both witness to be an addition to our knowledge. Now, what either of us may do with that knowledge is a different story. One thing is for sure, it ought to cause both of us to update our beliefs.

Thus knowledge (and information) bear on how we change our beliefs, and in that way—in a sense—become incorporated into our beliefs, but there is no rigid connection between the two concepts. Knowledge/information, as it is encoded in measurement outcomes, is a bridge to the external world that the quantum state has no right to be.

You may also recall another strange phrase I used in my note to Mermin: "the world (i.e., the detector clicks, whose meanings in the end are also set by subjective considerations)." This oddity is reflected in my definition of knowledge above: that is, I make a distinction between the raw stuff of the world that the measurement intervention brings about and the registration b in our noggins (as a flag for further actions in our role as agents). What I am thinking here is something roughly like the following. Take the famous white-on-black or black-on-white visual illusion that can be viewed either as a vase or as two faces facing each other. The raw stuff of the world may be compared to the ink and the paper giving the image. In order to say, however, that Mabuchi and I gain the same knowledge in viewing this we need the deeper cultural agreement that we will *both* call it a face or instead a vase when we see it.

Below I will put a glossary that tries to summarize where I have come so far in my attempts to make sense of quantum mechanics. Essentially, I'm expecting only the two terms above to be relevant to the fights we'll be having in writing our RMP article. But maybe it is nevertheless useful for me to lay my full set of language oddities on the table.

Lately, I've been jokingly calling my view (as it stands) the "sexual interpretation of quantum mechanics." (Most people turn red and become uncomfortable when I do that and explain why. I suspect the same will be true even in your reading of this note. So, brace yourselves.) The essential idea is that something new really does come into the world when two of its pieces are united. We capture the idea that something new really arises by saying that physical law cannot go there—that the individual outcome of a quantum measurement is random and lawless. The very fact that the consequence of the union is random signifies that there is more to the sum than is contained in the parts. But I promise you I won't reflect the licentious details of this view in the glossary below. I'll leave the missing terms to your imagination.

- ACT – The actual carrying out of a quantum measurement/INTERVENTION, after a DE-

CISION has been made by an AGENT to do so.

- AGENT – Any participant in the construction of a scientific theory. In older language, the observer.
- BELIEF – In the context of quantum discussions, a quantum state. Or one might say the quantum state is a compendium of BELIEFS.
- CONDITIONALIZING BELIEF – The rule one uses to update one's BELIEF consequent to the completion of a measurement INTERVENTION. In the language of Kraus and Preskill this would be called the "quantum operation" or "superoperator", respectively.
- CONSEQUENCE – Whatever it is that a measurement INTERVENTION elicits out of the world.
- DECISION – An AGENT, within his power can decide to perform one ACT or another upon the world. Just as physical law cannot impinge on what determines the random outcome of a quantum measurement, neither can it impinge on the mechanism behind an AGENT's decision.
- FACT – This is a word I do not like. One might have said that the outcomes of quantum measurement could be called facts just as well as CONSEQUENCES: But the word fact, to me, contains the connotation of a kind of permanence that I do not see in the quantum world. Facts are irreversible additions to the furniture of the world. But measurement INTERVENTIONS (and their CONSEQUENCES) can be reversed through the agency of a further outside intervener.
- INTERVENTION – The physical act that we call in older language the measurement of a POVM.
- KNOWLEDGE – One's mental representation of the obtained CONSEQUENCE of a given INTERVENTION into the world. Implicit in the use of this word, is that all communicating parties agree to the meaning of the given INTERVENTION, i.e., that it is this POVM rather than another.
- PROPERTY – A property is something possessed by a FACT. I don't like the word FACT.

04 September 2001, to Mermin, "Brilliance"

Merminition 23 :

> CAF Said: I am becoming more and more dissatisfied with the slogan "A quantum state is a state of knowledge, not a state of nature." ... I think we should get more into the habit of calling a quantum state a state of BELIEF.

Brilliant! All kinds of trouble would have been avoided.

You know, I'm not one to turn down a "Brilliant!" But your second sentence does clash a little with what you wrote on August 8:

Merminition 24 : *It seems to me Chris is getting much too subtle about this. I would talk about knowledge, not belief.*

All kinds of trouble WOULD HAVE BEEN avoided?

Speaking of brilliant—real brilliance this time—today I'm going into Munich to talk to Hans Briegel about his papers with Raussendorf. (I told you I would be in Munich, but I'm actually in the little village of Zorneding outside of Munich.) I think there's something very deep in them, if they hold up. You may recall I recommended them to you once.

I've got more things of a philosophic nature to write you, but I've just got to find some time to do it. I'll try to be back to the waves tomorrow.

04 September 2001, to Caves-Schack, "Objective Probability"

Cavesism 9: *I expect you to have a really hard time with this—please skip the lectures on my not being sufficiently Bayesian—but it is, in my present view, a necessary feature that expresses the tension that exists in the notion that maximal information is not complete. The state assignment can't be verified by examining the system, but it can be verified by examining the trail of evidence from which I acquired maximal information. If someone else finding that trail of evidence could say that he didn't have maximal information or that he had different maximal information, the notion that the information is maximal would be untenable, since apparently something further would be required to make it so. This seems like a natural for someone who takes seriously those quotes about the process of intersubjective coming to agreement. It grants to maximal information in quantum mechanics some, but not all of the properties of maximal information in a realistic world.*

It is hard for me to understand what that "trail of evidence" is a stand-in for if it is not a compendium of OBJECTIVE probabilities. You follow that trail, and you have NO CHOICE but to assign all the probabilities that the Gleason theorem gives (presumably if you are rational). So, pure quantum states give rise to "propensities" ... when those pure states are "right"? Is that what you are saying? (I said all of this, of course, in my original longer note, but it seems good to isolate it here.) Can you give me an operational definition of this notion of "propensity"?

And why can we toggle these propensities from a distance? Are you giving up on spacetime after all? Or is this a new way of applying the principal principle?

Now I really do have to join the family.

04 September 2001, to Caves-Schack, "Re: Fourth and Fifth Reading"

Cavesism 10: *The point of our conclusion is that the Dutch-book argument leads to a unique probability or density-operator assignment in the case and only in the case of maximal information. This is just an entirely different thing from using frequency data—or something else—to specify every component of a density operator.*

I still don't get it (though I've had a lot of wine by now). By hook or crook, I use the information available to me to assign a probability distribution over the outcomes of some informationally complete measurement. That assignment gives rise to a unique density operator.

I'm still having trouble seeing what is special about a "maximal information" assignment. I'm not lying; I'm not trying to cause trouble; I'm just not seeing it. (Think of me as the second referee of the paper. Would that be ethical?)

Good night!

05 September 2001, to Caves, "Noncontextuality Again (and Again)"

Cavesism 11: *The underlying structure is a specification of alternatives that can be grouped in various ways—these are the contexts—to make exhaustive sets. We are required to make noncontextual assignments; otherwise we are ignoring the fact that this specification doesn't distinguish an alternative in two different contexts. If it did, we would be dealing with a different specification. This is the perspective of my first point of view, which justifies noncontextual probability assignments in quantum mechanics from the Hilbert-space specification of alternatives.*

I still don't entirely get this. You say we are required to make noncontextual assignments, otherwise we would be ignoring the fact that the original groupings do not distinguish an alternative in two different contexts. But why could we not ignore it? Perhaps the underlying structure is there for an entirely different reason than something to do with probabilities? For some reason, this point of view is just not clicking for me.

Cavesism 12: *Your perspective is different. As I understood it, you think of an alternative in different contexts as a single alternative because it has the same probability in all contexts. But where did you get this equal probabilities? Surely they're not measured or determined or anything like that, since they are states of knowledge. In quantum mechanics you get that they're the same because the standard quantum rule says so, but this is using noncontextuality, not justifying it.*

I don't know what more to say on this. It means that identifying this consequence of this intervention with that consequence of that intervention is a SUBJECTIVE judgment. (I.e., that identifying this outcome of this measurement with that outcome of that measurement is a subjective judgment.)

Cavesism 13:

> CAF Said: What you say "needs" to be done, seems hopelessly underconstrained to me too. But I think you shouldn't view the problem that way. The point is, one simply has a way of clearing the air of the noncontextuality issue BEFORE getting down to the nitty-gritty of quantum mechanics. Noncontextuality is the base assumption upon which one plays a new game: What *physical* assumption makes it so that our instruments should correspond to POVMs and not some other mathematical structure?

I'm going to adopt your strategy, and just flatly say I don't get this.

Let me try again. Here is the game we should be playing. In the most general terms, a measurement is defined to be a group of elements (satisfying some given property) drawn from a set with a given structure. The individual elements correspond to the outcomes of the measurement. The question is, what should that structure be? What should that property be? What are the reasons for those choices? This much we will safely assume (for the reasons given above): The probabilities of the outcomes should depend only upon the individual elements, not the group.

That's all I'm saying. Here is an example of dumb theory.

> UNDERLYING STRUCTURE = one-d projectors onto a complex vector space.
> GROUPING PROPERTY = choose any set of projectors that project onto a complete set of linearly independent vectors.

Then the only probability assignment that can be given to the outcomes of such a notion of "measurement" is the uniform distribution.

So, we start over and say, "Maybe the grouping property ought to be that the projectors add up to the identity." Aha, that gives us quantum mechanics. But you see, there are any number of other combinations of structures and properties one might have played with. The question is, what is essential about the structure and grouping properties that we do use? By saying that we have cleared the air of noncontextuality, I simply mean that the existence of noncontextuality in the probability assignments should not be a question. It was settled before we ever started the game.

Cavesism 14: *Well, when I say "Hilbert-space structure of questions," I clearly don't mean only that there is a Hilbert space, but that the questions correspond to one-d projectors. That's why I add "of questions" to the phrase.*

No, what you mean precisely is: A "question" corresponds to a set of one-d projectors that sum up to the identity. So why don't you just say it in a precise way rather than a vague way? If I were uninitiated to quantum mechanics, I might have thought that you meant the dumb theory above. I'm serious about this.

05 September 2001, to Caves-Schack, "Unique Assignment"

Cavesism 15: *The point of our conclusion is that the Dutch-book argument leads to a unique probability or density-operator assignment in the case and only in the case of maximal information. This is just an entirely different thing from using frequency data—or something else—to specify every component of a density operator.*

I guess my trouble stems from one of the things I said in the long note announcing my worries about "Making Good Sense." There I said:

> I think the best we can say is: If Alice and Bob both believe adamantly that they know the answer to some potential measurement AND that measurement happens to be the same for both of them, then Dutch-book consistency and Gleason's theorem will enforce that they make the same probability assignments for all other measurements (i.e., that they assign the same quantum state). But said that way, I don't think any non-Bayesian will be particularly impressed: For they would say that all we have shown is "if two people know the same things, then they will know the same things." Woop-ti-do.
>
> The Bayesians among us will still have some room to be impressed: For it will not be a priori obvious to them that beliefs about one observable should have anything to do with beliefs about another. In particular, it might even surprise them that a common belief in certainty (for two observers) for any fixed observable should lead to equal probability assignments for all other observables.

What I am wondering is: What would impress a devout Bayesian (who is just learning quantum mechanics) about our argument? Thus, given what I said above, I wonder whether he would not be equally impressed by the following. By hook or crook, Alice and Bob individually come to their own subjective probability assignments for the various outcomes of a single informationally complete POVM. Then, because of Gleason, they will have to match in their subjective beliefs about the outcomes of all other measurements they might perform. That matching made no use

of the concept of maximal information. What does the maximal information case give us in shock value?

I think what you're going to say is that in the case of nonmaximal information, Alice and Bob may have come to different probability assignments for that informationally complete observable. And that they couldn't have done that if they had had "maximal information" in the first place. But as all this debate has already shown, I think I reject that position.

As my glossary from yesterday attests, I think what is going on with me is that I am becoming ever more uncomfortable with identifying quantum states with information, maximal or otherwise. Thus, instead of calling a pure state maximal information, I am becoming more inclined to something like:

$$\text{pure state} \;=\; \text{maximally tight belief (or judgment)}$$

or

$$\text{pure state} \;=\; \text{a nonrefinable belief (or judgment)}$$

or anything else more along those lines.

Leave the word information for what we gain when we see the outcomes of measurements. This entity we know two disparate observers should agree upon—by definition—if they are in free communication with each other. But the quantum state, on the inside of one's head, is a more personal state of affairs.

I'm going into Munich for much of the day (to visit with Briegel), so you may not hear from me again until tomorrow.

05 September 2001, to Caves-Schack, "Re: Identity Crisis"

I've now given your long note the fourth and fifth readings it deserves. (I'm holed away in an office near Briegel's.) And I'm not sure how to respond yet. I think I will await your responses to my other notes first. I think it is clear in the time that has elapsed since our first communication that we have moved further away from each other's position. Or I should say I've moved further away, from our original position.

Briegel just came; I'll be back.

05 September 2001, to Comer, "Malleable World"

Just a little quote I liked.

> Once you bring life into the world, you must protect it. We must protect it by changing the world.

From: Elie Wiesel (b. 1928), Rumanian-born U.S. writer. Interview in Writers at Work (Eighth Series, ed. by George Plimpton, 1988).

06 September 2001, to Caves-Schack, "The Underappreciated Point"

My fit of insomnia is running out, so I'm going to have to go back to bed soon.

Schackcosm 3 : *Please give me some feedback on these thoughts. I know that you wrote A LOT more, but I find it easier to go through your emails paragraph by paragraph.*

You know, of that LOT, most of it was written very carefully and very purposefully, so I do hope you will try to read it all with that in mind before trying to get me to readdress too much. (You can ignore the stuff on noncontextuality for now; that's not important for present issues.) Next week I won't have the leisure of writing too many notes, as I'll be out on the election trail trying to stump this Bayesian point of view ... and trying to make it LOOK consistent. (I.e., I'll be at the quantum foundations meeting in Ireland.)

But, before crawling back into bed, let me address your greatest FEAR:

Schackcosm 4: *The most important thing to remember is the limited scope of the paper. It tries to show that Bayesian probabilities do have a place in qm, not more. Remember that most physicists would reject this. We show that, contrary to conventional wisdom, subjective quantum probabilities are not arbitrary. Let me remind you that you agreed to a paper of this limited scope last year in Montreal. I do not want a paper that is significantly expanded.*

This must be addressing my single remark of September 3,

> Of course, as you know, I view this as a good opportunity (with page limitations no longer of great concern) to expand some points and give some more references. I think we can read the referee as agreeing with that.

because, in all my voluminous letters, I don't think I ever mentioned modifying the paper otherwise. Of course I think it is unhealthy that your first trip-up would be that point ... but we are all moved by different things. (I don't fathom your pressures, and you don't fathom mine.) In any case, let me say this for the record to try to clear the air:

> If given free reign—which I do not actually want—I think I could modify the paper to my own tolerance without changing its length at all, or, at most by a paragraph. I would be more than happy if you and Carl would just find a language I could agree to, and modify the paper accordingly. The phrase "more references" was a euphemism for citing more of Chris's papers.

Deep in my heart, I believe you guys fool yourselves in thinking that this paper will be more widely read simply by being short, but that is not the issue (and it has never been the issue). My passion is to get quantum mechanics straight: So, let's get it straight.

On the whole, in reading your two notes, I found your method of expression better fit to my present mentality than Carl's. Maybe I'll give more specific examples tomorrow. For the present, let me just mention two things.

Schackcosm 5: *It does indeed not follow from our Dutch book argument alone that two agents must agree on the maximal info they have. But suppose agent A has maximal information and agent B insists on assigning a pure state that is not consistent with A's information. A can then extract money from B. I don't think the symmetry of the situation is a problem here. From A's perspective, B is wrong, in the same sense of wrong as if A had a piece of classical knowledge that B chooses to ignore in a bet.*

Perhaps the greatest life change I have had is that I no longer like the phrase "maximal information" in this context. That little phrase carries with it an entire philosophy, and it is one that, to me, does not seem consistent with its roots and, more importantly, does not seem right.

The most one can say on Bayesian principles is that:

> From A's perspective, B is wrong. And from B's perspective, A is wrong.

If A and B can have two pure-state assignments, and the most one can say is the item above, then pure states should not be called "maximal information." They are maximal "something else", but it is not information. (In another note, I have outlined what I think that "something else" is.)

Carl thinks he can fix this by invoking a "trail of evidence" that uniquely fixes which of two pure states is actually the case. But let me juxtapose two of his paragraphs and then try to reemphasize the underappreciated point.

> Paragraph 1:
> In quantum theory maximal information also constitutes a belief, but we resist the notion that it corresponds to some objective reality out there. Why this resistance? Ultimately it's because the maximal information leads to a pure-state assignment that gives probabilities whose only reasonable interpretation is subjective. It is very important to remember that this is the primary motivation for much of what we do. Probabilities are the subjective language used to deal with situations of uncertainty, so wherever we find them, they must be subjective. The subjective view of pure quantum states gains additional support from the fact that a pure-state assignment can't be verified by consulting the system—the same can be said for a probability assignment—and the fact that a state assignment for a distant system changes when we obtain information about it without ever getting close to it—this also holds for correlated probability assignments.

and

> Paragraph 2:
> The state assignment can't be verified by examining the system, but it can be verified by examining the trail of evidence from which I acquired maximal information. If someone else finding that trail of evidence could say that he didn't have maximal information or that he had different maximal information, the notion that the information is maximal would be untenable, since apparently something further would be required to make it so. This seems like a natural for someone who takes seriously those quotes about the process of intersubjective coming to agreement. It grants to maximal information in quantum mechanics some, but not all of the properties of maximal information in a realistic world.

If you hold fast to the view that that trail of evidence must EXIST, then you hold fast to the view that quantum probabilities (in some cases) must be objective after all ... INDEPENDENTLY of the issue of intersubjective agreement. And that negates Paragraph 1.

I am now of the opinion that if we can just clear the air [I'm fond of that phrase] of this nonBayesian trapping from bygone times, we will finally be in a position for real progress. It is in Dutch-bookian type coherence (as a general principle) that one finds an objective statement in quantum mechanics; it is never in the quantum state itself, even when that state is pure. The OBJECTIVE statement is: All of you, each and every one of you, should manipulate your compendia of beliefs according to the rules of quantum mechanics if you wish to maximally avoid undesirable consequences in your gambles. The particular quantum states at any one time are just thin films of subjectivity floating on that wider sea of objectivity.

But, please, please do read the other notes carefully. I can only write a finite amount. I'll comment more particularly on your present notes tomorrow (i.e., today, after I get back up).

06 September 2001, to Caves-Schack, "Another Way"

> If you hold fast to the view that that trail of evidence must EXIST, then you hold fast to the view that quantum probabilities (in some cases) must be objective after all ... INDEPENDENTLY of the issue of intersubjective agreement. And that negates Paragraph 1.

Let me put it another way. By Carl's view, if trails of evidence MUST exist, then quantum states MUST exist, and the ghost of Bruno de Finetti should have stayed in the netherworld. For the probabilities derived from the quantum state will exist after all.

06 September 2001, to Caves-Schack, "Some Comments"

Now I return from an unrestful morning in bed.

Schackcosm 6: *I'd say "subjective but not arbitrary". Don't forget that the idea that quantum probabilities can at all be viewed as subjective is shocking to most readers.*

It won't be particularly shocking if, in the end, we reinstate their objectivity.

Schackcosm 7: *I guess you are right that we should be more explicit about "whose certainty". It is the bettor's certainty.*

The deeper issue is not that we *should* be more explicit about "whose certainty," but *why* we should.

Schackcosm 8: *You should leave Schopenhauer and your mother out of this discussion. The distinction between the cases of certainty (classical logic) and reasoning in the face of uncertainty (probability theory) is useful.*

I didn't understand this comment.

Schackcosm 9: *As I said in my previous message, two agents having conflicting certainties is a completely classical situation. If you accept classically that in this case, one of them must misread or ignore some of the available information, then the point of the paper is that the same classical argument gives you unique state assignment, even though states are Bayesian. This is a forceful conclusion.*

It is safer to have the wrong metaphysics in the classical case. This is because certainty (i.e., overpowering belief in the outcome) for one question means certainty for all questions. And that certainty can be verified or falsified in a single shot. So, one gets in the habit of thinking that the proposition (or its material counterpart, as instantiated in the world) was already there before the verification. One can accept that metaphysics or leave it, but it is usually more convenient to accept it. In the quantum case, however, if you assert that the proposition was already there (say, as uniquely specified by Carl's "trail of evidence") then you have to assert that all the rest of the quantum probabilities were already there too. That sounds an awful lot like objective, agent independent probabilities to me.

You can retreat to objective probabilities if you wish. But I say it is better to be creative with our metaphysics. JAW said it like this, "No question, no answer." And that distinction is rearing its head in this very problem.

Schackcosm 10 : *Making a pure-state assignment is an extreme statement. It entails the conviction that assigning a different state is equivalent to handing over money. It entails the conviction that the agent assigning the different state is wrong in this sense, in the sense of irrational behavior, not in the sense of not conforming to reality.*

It "entails the conviction." That is language I can accept. It is language I like. Trying to instate that way of saying things has been the whole point of my writing such detailed notes, especially the point about "not in the sense of not conforming to reality." But though you use it so nonchalantly now, it had no representation in our previous discussion, and it has no representation in our paper. At least looking at myself personally, I feel as if I have come through a phase transition.

Schackcosm 11 :

> CAF Said: But said that way, I don't think any non-Bayesian will be particularly impressed: For they would say that all we have shown is "if two people know the same things, then they will know the same things." Woop-ti-do.

Still quite impressive. A and B know the same certain thing. Hence they must assign the same subjective probabilities to all questions. Even subjective probabilities 0¡p¡1 are prescribed by this knowledge. The non-bayesian should be quite surprised and impressed by this.

I said non-Bayesian. Non-Bayesians do not accept subjective probabilities.

Schackcosm 12 :

> CAF Said: For I think in no way have we shown that when two observers make two pure-state assignments for a system, those pure states MUST be identical.

If they are not identical, each agent has perfect reason to assume that the other one is unreasonable.

I accept that. But the point has been, and remains, that that is the ONLY conclusion we can draw.

Schackcosm 13 : *There are dictatorial constraints only in the limiting case of maximal information.*

Unless all of this email has been a grave mistake on my part, I continue to not be able to accept this. The only argument we have at our beck and call is that Dutch Book + Gleason dictates what *I* must do in my head and what you must do in your head. It tells us each how to translate a *nonnumeric* belief (certainty) about the outcomes of a single question, to a *numeric* belief about the outcomes of all possible questions. Indeed, I will lighten up: For a raw BAYESIAN that must be quite an impressive conclusion. There must a good way to say that in the paper. [For the non-Bayesian however—one with no qualms about objective probabilities, one with no qualms about the objectivity of quantum states—I remain in my belief: it will strike him as little more than a tautology.]

But all of this does not lessen my debate with Brun, Finkelstein, and Mermin which started this whole affair. There is no a priori principle in the universe that will tell us that two quantum states OUGHT to have overlapping supports. The best one can say is that IF Alice and Bob have overlapping support, then (if they wish) they may be able to communicate the reasons for their beliefs and come to a more refined consensus. If they do not have overlap in their supports, then the only they can do to lessen their strife is consult the world.

The objectivity is not in the states, but in the state-space STRUCTURE and in the answers the world gives us upon our consultations. When one has gained the latter, one has gained information. But the quantum state before and after remains belief, pure state or not.

06 September 2001, to Caves-Schack, "Weak Point"

Cavesism 16: *I think we all agree that if states are Bayesian, then anyone can assign any state he pleases, including any pure state. He can be misled or tricked, or he can just be crazy, but this sort of freedom to assign any state is not of much interest for our paper. An objectivist will have no trouble agreeing that someone who is misled or irrational will use the "wrong" quantum state.*

Let me try to say it again. The main point is, in the quantum mechanical world, these "trails of evidence" you are thinking of in the back of your mind are NEVER enough to uniquely specify a quantum state. It has NOTHING to do with being misled or being irrational. Even the purest of states is thoroughly infused with belief from the get-go. That is what my note titled "Fw: Intersubjective Agreement" from September 3 is essentially about. So, this is not a case of measure zero, where the players are irrational or dumb to begin with ... so long as we take our own arguments about subjectivity seriously.

Cavesism 17: *What he wants to know is whether scientists acting like scientists, sharing all information in a spirit of genuine co-operation, mutual respect, and dedication to truth, can assign different pure states. And we show that scientists acting like scientists can't: sharing maximal information, they must make a unique quantum state assignment.*

If that is what he wants to know, then he is not going to find it from anyone's Dutch-book argument: our last one, or our slightly modified new one. Your point is a weak one. The Dutch-book game is an adversarial game. Anyone whose intention is to make his opponent go bankrupt is NOT going to share everything he knows with him. He will be silent and bet his money.

Cavesism 18: *What rescues this conclusion from trivia? First, it answers the question of why science doesn't go down the drain: subjective state assignments are constrained in the case of common maximal information.*

I don't believe the conclusion is trivial; I said this to Ruediger yesterday. But I also don't believe it has anything to do with rescuing science. Playing by the quantum rules ought to be enough.

Cavesism 19: *Second, it answers the question without referring to real, verifiable properties of the system in question. In a realistic world one might justify the agreement in the case of maximal information by saying that any disagreement can be resolved simply by looking at the system and seeing who's wrong.*

Disagreements in the quantum world are resolved also simply by looking at the system. Suppose you and I agree to everything in the world EXCEPT the quantum state for a given system. How do we resolve our dispute? We perform a maximally refined quantum measurement (a POVM with rank-one elements). We agree on the system's state thereafter. That is all that has ever been important in science anyway—that the world provides us with a way to COME INTO agreement for all FUTURE predictions. For god's sake, Albert Michelson did not believe that the speed of light could be a constant. But his tenure was not stripped away when he found a negative result. He revised his "impossible" belief and got over it.

ALL that one need to demand from a theory is that it provide a way for two agents to come to agreement for all FUTURE predictions. Quantum mechanics (surely) satisfies that. It has nothing to do with re-objectifying quantum probabilities, and I can't see that it has anything to do with this stuff we got in the habit of calling "maximal information."

06 September 2001, to Caves-Schack, "Re: The Well Appreciated Point"

Your notes are well-appreciated themselves; I am finding reading them productive. Unfortunately, I cannot reply in detail tonight, but I hope more will be waiting for me tomorrow morning.

Let me do say though that I think I addressed some of your points in the note I just sent to Carl (and CC'd to you). The main thing was this:

Schackcosm 14 : *What the Bayesian can say is: If A assigns a pure state, he knows with certainty that any other pure-state assignment is foolish (handing over lots of money).*

It is not a situation that can be resolved within science, by discussion or experiment or comparison of notes. Both A and B are certain there is nothing that could change their belief. For A, B could just as well reject all of quantum mechanics.

It is not a sin for A and B to disagree about the present. What would be a sin is if they could not come to agreement in the future. And quantum mechanics provides just such a mechanism. It is not true that experiment cannot change (absolutely firm) beliefs in the quantum world: quantum measurements are invasive, and thankfully so. Each measurement gives us the opportunity to throw away the past and start afresh.

07 September 2001, to Caves-Schack, "Email Not Received"

Ruediger wrote:

Schackcosm 15 : *But, as I argued in my last email, I think that a modified betting argument, now having A and B as adversaries (Carl thinks that this modified argument should not be called a Dutch book argument), shows that starting from two different pure states to come to a later agreement is not what science is about. A must dismiss B as a crackpot. This argument would be useful in the Peierls debate. But maybe not for our paper.*

I didn't receive such an email; can either of you send it again?

Indeed I would bet that science cannot be made in a (purely) adversarial environment. Science is about cooperation, trying our best to come to a consensus. (That is why I have not lost heart in writing all these ridiculous emails!) But, nevertheless, from time to time I do talk and try to come to consensus with people I deem crackpots. The point is, though someone may be adamantly wrong about ONE thing (say, a pure-state assignment from my perspective), it does not mean that he is adamantly wrong about ALL things. And therein lies a backdoor for a discussion with such a person.

The only thing that one has to trust in the making of science is that one's colleague is internally consistent. It is OK if he got SOME of the "facts" wrong (from my perspective), and that I got SOME of the "facts" wrong (from his perspective). It is enough that he is willing to join in with me in letting the world pull us together. I.e., that each of us is willing to participate in trying to convince the other that he is wrong by consulting the ultimate arbiter.

But, I'll write more later (in the context of your last two notes).

07 September 2001, to Caves-Schack, "(Backbreaking) Analysis"

Schackcosm 16 : *We need to find some common ground.*

Yes, that is true. And I think we already have some, maybe even a lot.

But, as I see it, there still remains a significant amount of trouble in the language we chose to use in the past ... and that is what is putting stoppers on our progress in the present.

Of course, I feel like I am repeating myself over and over, but let me go to your explicit "common ground" paragraph and try to lay out what I like and what I don't like about it. The thing that keeps me going is the hope that maybe this whole debate is a lot like beer: On their very first taste of it, most kids think it is a foul stuff. But after more and more of their friends offer it to them over time, it starts to become a pleasant diversion.

Schackcosm 17: *Assume A has information of the kind we call maximal, i.e., A knows that a measurement of a POVM containing the 1D projector P will give the outcome corresponding to P. Then assigning any state but P will be Dutch-book inconsistent. This will be A's inconsistency with her own belief. The beliefs of the bookie or of Nature do not matter. A knows that assigning any other state would make her accept a bet in which she (not Nature or the bookie) knows that she will lose for any outcome she believes is possible.*

Now assume B has access to the same piece of maximal information. B then knows that a measurement of a POVM containing the 1D projector P will give the outcome corresponding to P. Then assigning any state but P will be Dutch-book inconsistent. This will be B's inconsistency with his own belief.

Hence: Two agents having access to the same maximal information MUST assign the same state.

Here is how I would reword it to suit my present tastes.

> Assume A is absolutely sure that a measurement of a POVM containing the 1D projector P will give the outcome corresponding to P. Then assigning any state but P will be Dutch-book inconsistent. This will be A's inconsistency with her own belief. The beliefs of the bookie or of Nature do not matter. A knows that assigning any other state would make her accept a bet in which she (not Nature or the bookie) is absolutely sure she will lose for any outcome.
>
> Now assume that B is absolutely sure of the same thing, i.e., that a measurement of a POVM containing the 1D projector P will give the outcome corresponding to P. Then assigning any state but P will be Dutch-book inconsistent. This will be B's inconsistency with his own belief.
>
> Hence:
>
> > Two agents having the same absolutistic belief about the outcome of a measurement containing the projector P *MUST* assign the same quantum state.
>
> Or equivalently (but, to me, more forcefully):
>
> > Two agents having the same absolutistic belief about the outcome of a measurement contain the projector P *MUST* assign the same (subjective) probabilities to the outcomes of all measurements that can be contemplated.

What were the main substitutions? Essentially, they were simply:

"maximal information" $\longrightarrow$ "absolutely sure"
and
"maximal information" $\longrightarrow$ "absolutistic belief"

And the same substitutions count for the word "know."

To me, those simple substitutions completely change the metaphysical complexion of the statement. The statement goes no further than it has to go to make the quantum Dutch-book theorem stand its ground. Why go further?

Carl gave his reasons: To save science. But I do not see that is necessary in any way, and I do not personally believe that that method is on the right track.

What is wrong with the word "know"? To my Western-trained mind, it conveys the idea that there is something in the external world (the world outside of my head and beyond my control) and that my mind contains a mirror image of it. It conveys the idea that the outcome to the contemplated measurement already exists "out there" in some deterministic or fatalistic sense. It conveys the idea that I really need never have a look to see if the outcome is produced by my measurement: It's already there, and I know it; why waste time on a measurement? Notice that I let the word "knows" stand when it came to describing the very logic of the Dutch-book argument.

What is right about "absolutely sure" and "absolutistic belief" for me? They convey the feeling that what I have in hand is a belief, an extreme belief to be sure, but nonetheless a belief. That phrase never reaches out to the external world for its justification—or, at least it seems so to me.

What is wrong with "maximal information"? I think it screams out no more clearly than in your concluding statement above. In words that Gary Herling might use: The very phrase "SAME maximal information" is an abomination of the English language. In my mind, information is much like the word "know" (though a little looser in constitution). It too conveys the idea of a mind or a newspaper mirroring aspects of a preexisting reality. Besides that, the very fact that we have to go to the trouble to use the word "same" in conjunction with "maximal" conveys the feeling that the word "maximal information" was never appropriate in the first place. If information is some stuff we have gathered from the world AND it is maximal—the very most one can get—AND two agents really should be gathering up the same stuff, else one of them is wrong, THEN why do we have to go to the trouble of using the word "same"?

Well, we use it (i.e., "same") to keep ourselves from contradicting the belief that probabilities are subjective after all. Fine, so that is a good reason to keep the word "same". BUT, it is a bad reason to keep the phrase "maximal information" to merely convey the concept that one is "absolutely sure" of the consequence of some action that one might take (i.e., the measurement being contemplated here).

Schackcosm 18:

> CAF Said: Suppose you and I agree to everything in the world EXCEPT the quantum state for a given system.

It follows from what I just wrote that this situation cannot arise in the case of maximal information.

To me, that statement is a *non sequitur*. I cannot find any *logic* to bring it about. And I say that especially if you can agree to the validity of my attempt at expressing a "common ground" for us in the highlighted paragraphs above ... no matter how pedantic you think my actual phrasing is.

If the quantum state is not uniquely declared by some reality, then there is nothing to stop us from agreeing on some aspects of the world and disagreeing on others.

Please read the note I wrote Mermin titled "Intersubjective Agreement" again. If you and I (in the presence of each other) perform a given POVM consisting of rank-one projectors on a system, then you might say that we will agree on the system's state thereafter. But that requires the assumption that we BELIEVE the same quantum operation (for updating our states) will be

associated with that measurement. If we don't agree on that at the outset, then we will come to conclusion of two different pure states for the system after the measurement is completed.

You say, well a quantum operation is surely an element of reality: It is either right or it is wrong. (Carl would say that in any case; I wouldn't.) But suppose it is so—I will relax my debate with Carl for the moment. How would we know which quantum operation we had? We would have to have prepared a load of quantum states beforehand to map which quantum operation is "really" there. But then we would have had to agree on our cache of exploratory quantum states in the first place. How did we get to that stage of agreement, I ask you? And, on I will do the same, ad infinitum.

The point is, in a world where our only exploratory tools are quantum states and quantum measurement outcomes, we can never terminate the chain. This is one aspect of what I meant yesterday when I said that quantum states are infused with beliefs from the get-go.

Quantum measurement outcomes alone will never, ever be enough to uniquely determine a quantum state. One has to have some further a priori information or beliefs to do that. You can play the game—as Carl wants to—that that a priori information is the world's Hamiltonian. But then you will be about as stuck as Kant was with his transcendental idealism: you will still have to start off with agents of some initial common belief before they will ever be able to come to agreement about the Hamiltonian's form. And how are we poor finite beings to ever get to such a starting stage?

I say simply: throw out any trappings that a quantum state can ever be objective.

Schackcosm 19 : *Is there ever maximal information? Yes. You give us an example where C and F both obtain the same maximal information about the system:*

> CAF Said: We perform a maximally refined quantum measurement (a POVM with rank-one elements). We agree on the system's state thereafter.

In the sense that you want maximal information, i.e., something beyond absolutistic belief, I would therefore say that there is never any. My example required that the two agents share an almost strangulating amount of common belief. See discussion in previous section.

Schackcosm 20 : *As Carl writes, if C and F have the same maximal information, they must assign the same state. This is an important situation, as scientists share the information they have.*

If we can get past the language, I will (clearly) agree that this is an important situation. Scientists share the data AND the beliefs (interpretations, machine designs, etc.) they have.

Schackcosm 21 :

> CAF Said: The Dutch-book game is an adversarial game. Anyone whose intention is to make his opponent go bankrupt is NOT going to share everything he knows with him. He will be silent and bet his money.

No. The Dutch book game is about ONE agent's consistency.

Yes, you are absolutely right. And I apologize for throwing in extra junk that is not relevant. But the only point I really meant was that there is nothing in the Dutch-book set-up that forces the bettor and the bookie to share their information. That is an extra requirement if you want it. But it is a requirement that seems to me almost to give up the whole spirit of the Dutch-book situation: it involves no communication beyond the numbers p and x.

Schackcosm 22 : *That agents having access to the same maximal information must assign the same state is all we need for our paper.*

It seems to me, the only thing we need for our paper is the "common ground" statement I made above. I would not think that I need to say it again, but I'm getting pretty fearful of the phrase "maximal information." At the very least, I would like to start using it in a more limited sense or in a more limited way. Or perhaps in a vague enough way (for the present project) that I can worm out of it when I want to write future papers of my own without you two. (But this issue is likely to haunt us all the way through to the end of the RMP deal.)

All this email is starting to exhaust me, and it has certainly kept me from making this trip to Munich even resemble a vacation. I would like to draw it to a reasonable end soon. (But I do understand that that will require flexibility on all our parts, even me.) This may help my samizdat production, but it no longer feels like it is helping the rest of my life. I feel like I have hold of some important points that we were just too much in the "classical" tradition to recognize before. If we ultimately disagree, then we'll just have to do that, but I would rather not end up in that state of affairs.

If you have to make choices on what to do with your own time, please comment on my note titled "Note on Terminology" sooner rather than later. I fly out for Ireland Sunday morning.

08 September 2001, to Caves-Schack, "Negotiation and Compromise"

Let me tell you a little story I dreamed up while driving through the Austrian countryside today. It is based on one of the most annoying realities of my life: there are times when Kiki and I just cannot come to agreement. If I can use Ruediger's words, there are times when I just think:

Schackcosm 23 : *A is a physicist who would bet his career on his state assignment. If he says B is wrong he means this in an absolute, very strong sense. He has examined all the evidence, and there is no doubt left. He is certain that B is missing some evidence. The Dutch book argument shows that A is certain that B's position is equivalent to handing over B's entire fortune. Wrong implies foolish, deluded. For A, B is a crackpot, and the circumstance that B is certain that A is wrong reinforces this position.*

But there are realities: Kiki and I are married; we share a bank account. And here and there, Kiki consorts with the Dutch.

What I am leading up to is that I think there is a place for Ben Schumacher's observation about a three-person Dutch book in our ongoing debate.

You two want to believe that there are god-given constraints on how much two people can disagree. I say there aren't. It'll be a miracle if we ever come to some consensus on this. But I have never said that there ought not to be reasons that two people might want to come to agreement ... EVEN when they have differing but, nevertheless, "maximal information." (I use the phrase "maximal information" despite my dislike for it in order to be sarcastic and to underline a further difficulty with the term in a moment.)

I am internally consistent; there's no Dutch bookie who can take me to the cleaners. And despite my feelings for Kiki's complete foolishness, I feel that she is internally consistent; there's no Dutch bookie who can take her to the cleaners (as far as she is concerned). But we will be in deep trouble if that Dutch bookie approaches us separately. (Being married, we report all our beliefs to each other.)

I can see two outs to this problem. The first one—which is less interesting—is that we make an effort to come to agreement by consulting the world. We make a measurement, and thereby, through its invasiveness, force the quantum system into a state we can agree upon. (Assuming, as I keep harping on, we can agree on the quantum operation associated with our measuring device.) But what if we have no access to the system of interest? What are we to do then?

I think we would have no choice but to, each of us, back off in the firmness of our beliefs. That is, we should agree upon a density operator that contains in its support both of our earlier ascriptions. We both give up some of our certainty is this process, but the upshot is that we no longer have probability one of becoming bankrupt.

So think about this: Two agents start out saying that they are absolute in their convictions about a some quantum measurement. But then the reality of their partner's stubbornness hits them, and the only thing they can do is back off.

You continue to want to call a quantum state information. But for the present case, again, it seems the term "information" is stretched beyond common usage by the factors people must sometimes take into account in coming to their assignments. From my point of view, Kiki's foolish quantum state assignment is not information about the physical world at all. It is completely wrong, with no reflection in the world as far as I am concerned. Nevertheless, I had better take it into account in making my bets if I don't want to lose our whole joint bank account.

The point I take home is that is sometimes better to negotiate and compromise even when one has "maximal information."

10 September 2001, to Caves-Schack, "Short Reply"

Just a very short reply to your latest posting. But first let me say something about this:

Schackcosm 24: *I hope this email establishes a little more common ground. I would like to start, next week, on revising the paper in the view of this discussion and the referee's comments.*

Yes, do it. With the draft in hand, I'll be better able to see which statements make me feel like a liar and which do not. (Hopefully most of them won't.) And then we can be done with this, and then you can finally stop saying to me, "But not for this paper."

Schackcosm 25:

> CAF Said: What is wrong with the word "know"? To my Western-trained mind, it conveys the idea that there is something in the external world (the world outside of my head and beyond my control) and that my mind contains a mirror image of it. It conveys the idea that the outcome to the contemplated measurement already exists "out there" in some deterministic or fatalistic sense. It conveys the idea that I really need never have a look to see if the outcome is produced by my measurement: It's already there, and I know it; why waste time on a measurement? Notice that I let the word "knows" stand when it came to describing the very logic of the Dutch-book argument.

I do not think any of this is clear. I "know" something means I have a firm belief in it. I don't think more is implied. The question of whether something corresponding to the knowledge (or the belief) exists out there is entirely separate from these wording issues.

Mermin once wrote me this:

> It seems to me Chris is getting much too subtle about this. I would talk about knowledge, not belief. I take "knowledge" to mean simply "true belief", a definition that as I remember goes all the way back to Plato and can be made unproblematic even in the quantum context. [Only a postmodernist would sneer at my saying this.]

And then this:

> C'mon, don't drag your heels. QM is sometimes capable of assigning probability 0 to certain outcomes. For those one doesn't have to argue about whether probability has to do with ensembles or degrees of belief or anything else. "True belief" seemed a good term to describe such outcomes, and then I remembered that Plato (I think) had used the same term (in contrast, as I remember, to "opinion".)

William James writes it like this:

> The popular notion is that a true idea must copy its reality.

The *Encyclopedia Britannica* says this (mainly the last sentences are relevant):

> In general, the philosophical tradition from the Greeks to the present has focused on the kind of knowledge expressed when it is said that someone knows that such and such is the case, e.g., that A knows that snow is white. This sort of knowledge, called propositional knowledge, raises the classical epistemological questions about the truth or falsity of the asserted claim, the evidence for it, and a host of other problems. Among them is the much debated issue of what kind of thing is known when one knows that p, i.e., what counts as a substitution instance of p. The list of such candidates includes beliefs, propositions, statements, sentences, and utterances of sentences. Each has or has had its proponents, and the arguments pro and con are too subtle to be explored here. Two things should, however, be noted in this connection: first, that the issue is closely related to the problem of universals (i.e., whether what is known to be true is an abstract entity, such as a proposition, or whether it is a linguistic expression, such as a sentence or a sentence-token) and, *second, that it is agreed by all sides that one cannot have knowledge, in this sense of "knowledge," of that which is not true. One of the necessary conditions for saying that A knows that p is that p must be true, and this condition can therefore be regarded as one of the main elements in any accurate characterization of knowledge.*

I put stars beside the hot stuff. [It is italicized in this samizdat version.] And I'd send you more quotes if I could, but I'm a gazillion miles away from home in a piss-poor dormitory room (with no library).

Schackcosm 26 : *There must be a misunderstanding. If you and I agree about everything in the world, we also agree on the certainty of a particular measurement outcome. We can not then disagree on the state to assign, because at least one of us would be internally inconsistent.*

You must mean that we agree on everything except the state AND the measurement.

Yes.

Schackcosm 27 :

> CAF Said: I now think it is much better to reserve the word KNOWLEDGE solely for the outcomes of quantum measurements once they become part of the mental makeup of an agent interested in them. I walk into Mabuchi's lab, and to the extent that he and I agree that he is performing some POVM (denoted by a set of positive operators $\{E_b\}$), it seems to me valid to call the outcome b we both witness to be an addition to our knowledge. Now, what either of us may do with that knowledge is a different story. One thing is for sure, it ought to cause both of us to update our beliefs.

This is a very difficult debate. I am not sure I understand why you draw the line between knowledge and belief exactly where you do. Why give belief in b a special role? I need to think much longer about this.

Because I had assumed that Mabuchi and I had agreed to a fixed "random variable." With respect to that prior assumption the thing we gain is "information" in the standard Shannon sense. I'm willing to call that knowledge that we did not have before. There is nothing personalistic about it; we both have gained the same thing.

I went to dinner with Jon Barret, Harvey Brown, Matthew Donald, and David Wallace. Getting Bayesian ideas across to them is going to be a hard sell. But the most amazing thing is that each and every one of them was familiar with the Dutch book argument! (Now, how could anyone know it, and yet it not change their life?)

15 September 2001, to Schack, "Re: question on the manuscript"

Just send me the draft when you're pleased with it. And I'll tell you then what I can tolerate.

Schackcosm 28 : *But as your (ugly) term "absolutistic belief" shows, they should be extremely careful before making pure-state assignments.*

I am not committed to that term, like you guys are to the (inaccurate) term "maximal information." It was just the best I could come with at the time. Change it if you wish. The most important thing is that whatever you substitute for it should carry no flavor of a "reflection theory" of knowledge.

You won't buy it, and you'll think I'm just saying this out of less than pure reasons, but you would be surprised at how many people have now encouraged me to call the quantum state a "belief" rather than "knowledge." (Four.)

I'd like to come to SB, but we'll have to see how things play out with Kiki.

16 September 2001, to Mermin, "For Worse"

Merminition 25 : *Well for better or for worse Todd, Jerry, and I came up with something we could all agree on.*

Well, I managed to download your paper this morning after all. I ended up in the airport for over four hours. Now, I'm somewhere over the Atlantic.

I hate it, of course. It's built around the same stubborn disregard for the issues involved in trying to give the quantum state a NONontological status that all the BFM emails were.

To use a term I picked up from a philosopher this week in Ireland, the proposed outcome of a single quantum measurement simply cannot be a "candidate for knowledge." It's source must be considered ineffable, else the wrath of Gleason's theorem would strip away the nonontological status that was being sought for the quantum state in the first place.

I can guess from the format of the paper that it is being submitted to PRL. If I were being mean, I might say, "Good. That'll give it a higher chance of being rejected." But I'm not mean. If I were the referee I might even accept it. But, ONLY if the authors use the extra column at their disposal to explain why their statement of the problem is not a deep endorsement of the Everettista manifesto. (Or, if it is, then just come out and say it.) What else is Zeno but a baroque name for the quantum state of the universe (mixed or not) or a candidate for the one that really gets it "right"?

On a personal note: If you're going to cite a personal communication that appears in the samizdat, why don't you cite the samizdat? (I'm always looking for further ways to draw people into it.)

Best regards (in continuing disagreement),

16 September 2001, to Mermin, "Persnickety Business"

Since my last note, I've looked over your paper again. I must say I found the sentence,

> It is surely a significant feature of the theory that consideration of impossible outcomes and very little else, without any invocation of "the uncertainty principle" or "maximal information", to the fact that pure state assignments must be unique, as well as the more general constraint on mixed-state assignments.

a bit snobbish. Especially since you gave zero devotion to the issue of whether a quantum state could be "true" or not. If they can be "true," then why don't you just take them to be a bit of material reality and forget all this crap about knowledge? If they can't be, then what do you mean by the word "impossible"?

I still hate the paper, probably more so now.

A little further over the Atlantic (and wishing I were closer to Greenland),

16 September 2001, to Baker, "Arlo and Arlo"

As I just wrote another David—i.e., Mermin instead of Baker—right now I'm somewhere over the Atlantic, wishing I were a little closer to Greenland. I'm on my way back home from Northern Ireland. What a week it has been. I've never been so homesick for America before.

> Good mornin' America, how are you?
> Don't you know me, I'm your native son.

I think I might just kiss the ground if I really do land at JFK.

Once upon a time, we could actually see the twin towers from Morristown on a clear day. They pierced the horizon like nothing else in Manhattan. Hearing Emma's voice on the phone the other night was one of the toughest things. She said, "Hi Daddy-o," and then told me all about her first day in playschool. It's not easy to juxtapose that with the death of 5000 people and the many more deaths that may come in the near future.

Life is an essence in this universe. Creative, productive, reproductive life: I'm not one of those scientists who think it is just an illusion, an epiphenomena rolling on top of a clockwork or a dice-rolling world. It is something in itself. It was latent there all along, long before the trilobites had found their place, and it is our task to see that it does not return back to that latency.

Very sad times.

16 September 2001, to Renes, "Registered Complaints"

Good to hear from you; thanks for the note. I'm flying back from Ireland as I write this to you. I think the conference was successful in many ways. While there, I got to know Marcus Appleby, and he is a really good guy. He's very Jamesian in his perspective, very clear in thought, and a real seeker of the truth. I think I'm going to work to get him to Montréal next year (even though he does not practice quantum information).

I agree with your assessment of the BFM paper. I've already registered my complaint with Mermin. But you know, it's not like there are no pansies even in the midst of our own little group: Carl and Ruediger have had no great fortitude when it came to this issue, no courage of their convictions.

I'll place my small samizdat on the subject below in case you're curious.

I hope your transition back to New Mexico has gone well. We'll miss you at Bell Labs.

16 September 2001, to Schack, "Nerve Therapy"

I'm in the last 1.75 hours of my flight to New York, and I'm doing everything I can to keep from getting too nervous. Let me finally turn my sights to your note from a week ago.

Schackcosm 29 : *I was referring to the argument that, from A's perspective, if B assigned a different pure state, that would force B to accept a bet offered by A that amounts to handing over money to A. Therefore B is a crackpot from A's perspective (but not from his own, of course). This is interesting because it is the very fact that B has maximal information that forces him to accept A's bet: B "knows" A cannot know anything that would give her an unfair advantage in the bet. I still think this is a strong argument. Scientists would try to overcome this critical situation by questioning the world, i.e., by making measurements, but NOT on the system in question, but on everything else, e.g., Hideo's optical table. After all they both are convinced that the world is disentangled from the system in question.*

Personally, I only find this restriction to not touching the system of interest (by you and BFM) ad hoc. Since when in science do we restrict the experimentalists to never touch the objects of their interest? Besides, I think you may be missing one of the deepest facts of the quantum world: it has the power to cause agents to agree *henceforth*, no matter how adamant their previous beliefs ... AND EVEN without the underlying reality of preexisting measurement outcomes that the classical world had. With each quantum measurement, a part of the world starts afresh.

Schackcosm 30 :

> CAF Said: But there are realities: Kiki and I are married; we share a bank account. And here and there, Kiki consorts with the Dutch.
>
> What I am leading up to is that I think there is a place for Ben Schumacher's observation about a three-person Dutch book in our ongoing debate.

I cannot quite see where this situation offers any specifically quantum insight. In any case, it is completely different from the "adversarial game" I am talking about.

I'm not sure that it does lead to any specific quantum insight. But then again, I had never seen you and Carl be so adamant about taking probabilities to be objective features of nature before. (Don't do it! I know you're going to say "NO, NO, NO, we have never wanted objective probabilities!" But all the pieces of evidence—to me—point to the contrary.)

What it does show—and that is maybe where quantum theory comes in a little bit—is that there are times when no matter how adamantly I believe something, I should bet according to odds different from my belief. I.e., there are internal beliefs, and external "beliefs."

To that extent, I am starting to wonder if "belief" is even an appropriate word for capturing the essence of a quantum state. Instead it is starting to seem to me that it may more appropriately be described as a negotiated signifier to external action. I.e., in some cases, it signifies the betting strategies that a community of scientists can agree upon, even in the case of more refined and divergent beliefs.

Whose knowledge, Mermin asked? Sometimes, mine. Sometimes, yours. Sometimes it's more a matter of the policy we have been able to negotiate.

I do get a little worried in saying all of the above, in that, maybe I have not taken Savage's and Bernardo and Smith's views of probability into adequate account: namely, that probabilities and utilities spring into existence at the same time, and are a little inseparable. But presently I don't see how that fits this problem. Nor does it extinguish the problem that sometimes it really is in people's best interest to lie about what they believe.

Schackcosm 31 :

> CAF Said: I am internally consistent; there's no Dutch bookie who can take me to the cleaners. And despite my feelings for Kiki's complete foolishness, I feel that she is internally consistent; there's no Dutch bookie who can take her to the cleaners (as far as she is concerned).

I strongly believe that here marriage is fundamentally different from science.

Perhaps. But the point is, sometimes we bet according to situations that have nothing to do with our beliefs. We bet so as to obtain the best common good. And with that we do return to some aspects of science (and politics).

Schackcosm 32 :

> CAF Said: But we will be in deep trouble if that Dutch bookie approaches us separately. (Being married, we report all our beliefs to each other.)

This deep trouble arises even if both you and Kiki assign mixed states with the same support, I believe (but I need to think more about this). I am trying to make the point that pure-state assignments are different.

That is true. I chose the particular example of pure states to be as dramatic as possible ... in an attempt to hit you and Carl in a point where you were being the most hardheaded.

Schackcosm 33 :

> CAF Said: I can see two outs to this problem. The first one—which is less interesting—is that we make an effort to come to agreement by consulting the world. We make a measurement, and thereby, through its invasiveness, force the quantum system into a state we can agree upon. (Assuming, as I keep harping on, we can agree on the quantum operation associated with our measuring device.) But what if we have no access to the system of interest? What are we to do then?

Well, I gave my answer above. We question everything in the world BUT the system of interest. Using quantum measurements, of course.

Ad hoc.

Schackcosm 34:

> CAF Said: I think we would have no choice but to, each of us, back off in the firmness of our beliefs. That is, we should agree upon a density operator that contains in its support both of our earlier ascriptions.

Yes, the same is true for two reasonable scientists. But as your (ugly) term "absolutistic belief" shows, they should be extremely careful before making pure-state assignments. Your way out is similar to having second thoughts after placing the bet.

No. My way out is to point out that there is, after all, a distinction between an (internal) belief and its manifestation as an outward bet. The belief comes before the bet.

OK, I've just passed Boston, and in the mean time, got another meal in my stomach. For the obvious reason, getting past Boston carries some symbolic significance. Let us believe in symbols.

18 September 2001, to Peres, "Hi Back"

Good to hear from. Yes, I am home. I was delayed by a day, spent a night in London, had to return to JFK rather than Newark, and had the most tense flight of my life, but I am home. I held and held and held on to Emma when I saw her waiting for me.

I have been in a huge email debate with David Mermin about his two latest papers. Consequently, I also got into a huge debate with Caves and Schack, though they are (slightly) more sensible. (The whole thing has been quite taxing.) John Wheeler once told us a story of a condemned man who, while waiting for the firing squad, calmed his nerves by contemplating Hamilton's beautiful equations. I learned a lesson from this: When I was worried about the safety of my flight, I calmed my nerves by continuing my email debate with Mermin, Caves, and Schack!

I hope you and Aviva are able to get home on schedule. I suspect your grandchildren are waiting for your return with excited eyes.

18 September 2001, to Mermin, "Goodbye"

I just don't know what more to say. Clearly there is something that is keeping us from communicating. I write more and more, and you just don't get my points. It's probably better if I just write less and less.

I'll record the way Caves put the best part of your paper the other day and just leave it at that. I've never said anything different, but maybe if you hear it from another voice something will click.

It seems to me that there are various kinds of coming to agreement or inability to do so. Much of this is a recapitulation of what we have already discussed. It is phrased in a way that is supposed to avoid the notion that beliefs incorporated in a state assignment require the world to do things. [...]

1. BFM consistency (a multi-party condition): If the parties share their information, it is not ruled out that they can come to agreement on a common state assignment.

If they share information, they must rule out all vectors in the subspace generated by the union of their null subspaces, and they must assign nonzero probability to all vectors in the intersection of their supports. They have the possibility of making a joint state assignment if and only if the intersection of their supports is of dimension 1 or more.

Notice that we don't have to say here, as BFM do, that an outcome assigned probability zero by any party definitely cannot occur.

But let me say one last thing before I give up on this conversation. In your last note you wrote:

Merminition 26 : *We've sent the hateful paper to Phys Rev A as an ordinary low-grade submission. My concern is that a referee will condemn it as trivial or well-known, not outrageous or Everittistan.*

You have probably summed up the referee's response correctly, but it doesn't make it so. Nor does it make you paper non-Everettistan after all. The whole point can be summed up with the following part of your paper.

A necessary condition for compatibility

Suppose Alice, Bob, Carol, ... describe a system with density matrices $\rho_a, \rho_b, \rho_c, \ldots$. Each of their different density matrix assignments incorporates some subset of a valid body of currently relevant information about the system, all of which could, in principle, be known by a particularly well-informed Zeno.

This is an edict you are placing ON TOP OF quantum mechanics. It is a construction of your own doing. I find it in no axiom system I have ever looked at. It IS a good half of the Everettista starting point.

You might say the other half of Everett is that he takes the quantum state explicitly to be an ontological entity. But you've got that too. By imagining that facts of the world (measurement outcomes) uniquely determine a quantum state, what else could a quantum state be but a stand-in for those real things, those FACTS (the stuff that one absorbs into one's mind and calls information)? You simply cannot say that quantum states MUST be consistent, if you want them to retain a nonontological status. The second you say they MUST be, then you have given up the game: states are then properties of the world, and this question of consistency was a waste of time from the beginning.

With this, I end the conversation: I give up.

19 September 2001, to Caves-Schack, "Morning Coffee"

Schackcosm 35 : *Yes, and this is great progress!*

Well, the note really is some progress. Let me just accept the gifts the Lord has given me.

I have two questions.

1) What does Peierls' consistency have to do with coming to agreement? I don't see it yet. It seems to me to be more a technical definition for the word "wrong" than anything else. Using it, I can say when you are "wrong" with respect to my beliefs, but that seems about it.

2) This thing you call fuchsian consistency. You bring to the surface, a point I've always been a little worried about. Why refuse to consider outcomes lying in the null subspace of a density operator? Let me give an example:

> Chris: I say the state is spin-up in the z-direction.
> Carl: I say it is spin-down.
> Chris: Spin-up!
> Carl: Spin-down!!
> Chris: OK, let's test it. Here is an ideal z-spin-measuring device. It's Kraus operation is just a von Neumann collapse. Do you agree?
> Carl: Yes, I agree.
> Chris: Are you sure you agree? We need some agreement or this won't be a meaningful game.
> Carl: Yes, I agree and am growing impatient.
> Hideo: Spin measurement done. It's spin-down.
> Chris: Damn, I honestly believed with all my heart that it'd be spin-up!
> Carl: Well, I wouldn't have said it if I didn't know it.
> Chris: OK, from here out, I say the spin is spin-down. I was a fool before, but at least we agree now.

Your worry seems to be that the formal apparatus of quantum mechanics will not allow me to propagate my belief. According to the Kraus rules, my new state should be the zero operator renormalized by a probability zero: i.e., it is an undefined object. Nevertheless, life goes on, just as the story above indicates. This, at least, does have a formal counterpart in that, if I back off in my belief even the slightest amount to any mixed state (not on the boundary), then the Kraus operation will take us both to the same place, to the same one-d projector. Consequently, the limit exists even if the actual point is undefined.

So, in saying this, I don't think one has to do anything fancy with a measurement to get this thing you call fuchsian consistency off the ground. No intricate constructions of states seem to be needed. Choose ANY (mutually agreed upon) von Neumann measurement you like, and Carl and Chris will have to agree at the end of the day. One of them will feel foolish for having been so wrong, but as long as he is rational, agreement can still be had. As Ruediger said, the idea is that this is a property of quantum mechanics (in its capacity as a law of thought), not of the initial states.

Alternatively, one can turn the problem around, as I tried to do in my August 7 note to BFM: suppose Carl and Chris agree to a *particular* Krausian state change rule (one not even closely resembling a von Neumann collapse). Now one can ask, under what conditions on the initial states will C and C be left with more agreement at the end of the day than they started with? But that is a different problem than this f-consistency condition you speak of.

19 September 2001, to Mermin, "Tentative Hello"

Merminition 27 : *I may take a shot at rewriting our argument, since I'd like to produce something for the Växjö proceedings. At that point I may try it out on you, but if you think it's likely to induce nausea, you don't have to respond.*

I owe you many things. If you want me to look at it, I will look at it, and I will try to respond (i.e., give you constructive criticism) in a civil tone.

19 September 2001, to Caves-Mermin-Schack, "Practical Art"

I suspect all of you have seen, in one form or another, the optical illusion where, looking at the drawing in one way, it appears to be a beautiful young woman. But looking at it in another way, it looks to be an old hag. Here's a true story about that. The first time I ever saw such a thing was in Roger Penrose's book *The Emperor's New Mind* in 1989 or so. Below the picture was a caption, explaining just exactly what I ought to see: alternatively, a beautiful woman and a old hag. But the oddity was, I could only see the beautiful woman. As much as I tried to find a hag in the picture, I couldn't do it. That fascinated me—so much so that, from time to time throughout the next year, I would pull the book off the shelf and try to find the old hag. I never did see her until one day I was searching through the book for something completely different and happened to come across the page. In fact, that may have even been two years after my initial encounter with the picture.

I think it is safe to assume from this that if Penrose had never pointed out to me that there ought to be an old hag there, I simply may have never seen her.

Let me apply this piece of art to a question on my mind. Let us imagine that some perpetrator has committed a dastardly deed to my family and left a note promising that before the end of her life she would do it again. Beyond this I have no clue of the identity of the perpetrator (or of her cause) except a sketch made of her by an eyewitness, who has now also disappeared from the scene. For intentional completeness, let us suppose there is simply no further evidence that I can lay my hands on—there is no more "currently relevant information," no deeper "trail of evidence." The carbon atoms laid out on the two pieces of paper—the note and the drawing—are the only links I have to the cause of the crime. The trouble is, in the case of the drawing, the carbon atoms sketch out the shape of Penrose's beauty-hag.

Now suppose an insurance salesman appears on the scene and is willing to sell me insurance against a recurrence of the crime precisely 30 years hence at such and such a premium. Do I buy it? That, of course, will depend upon the probability I ascribe to the crime's recurrence in precisely 30 years. What probability do I ascribe?

It will depend upon what I see in the picture. If I "see" the old hag, it will be one thing. If I "see" the beautiful woman it will be another. If I "see" there is an ambiguity, it will be still a third. But one thing is sure from this example, what I "see" is not completely dependent upon the pattern laid out by the carbon atoms. Part of what I "see" is dependent upon psychological factors that I *myself* may simply never have access to—things deep within my head, things determined when I was first toilet trained, things to do with the self-referential nature of consciousness, things that physics just ought to (and maybe even has to) leave alone. My beliefs (my probability assignments) after seeing the sketch are determined in part by something objective in the world external to me, but also in part by my previous, purely subjective beliefs. This example ought to make it clear that I cannot put all the weight of my posterior beliefs on the external world.

Here's my question: Where do you draw the line? More importantly, how do you draw the line? What is it in your view that gives the clicks of a quantum mechanical measurement a status that goes deeper than the example above? What makes the phrase "currently relevant information" more decisive in the quantum mechanical case? How does that "trail of evidence" have more power to eradicate my subjective pre-beliefs than any other? So much so, in fact, that one gets the feeling that we might even view the quantum state as determined solely by such a trail? You say,

"Well, there the information conforms to a physical theory." But in what way does that change things? Can you pinpoint it? Can you make that declaration meaningful? Indeed, can you pinpoint what about the example above is not quantum mechanical? It had atoms, it had systems, it had probabilities.

I said I would not talk to David anymore, and here I am talking to him. I guess that just goes to show that that Mermin can steal your heart ... as can all of you.

20 September 2001, to Caves-Schack, "The Stopgap"

Attached is the latest draft. You will find two (and only two) changes.

1. I added the word "the" in front of "records."

2. I added two citations to myself. One at the beginning, one at the end. If you don't like where I put them, that's fine—feel free to change their positions—but I would like to have them somewhere in the paper.

Of course I am still troubled by the use of the phrases "maximal information" and "with certainty". I continue to think that they convey an image that is better left unconveyed, but I'm the odd man out here so I'll shut up. The main thing is that you left me enough wiggle room that I can defend my view of the quantum state when I need to. E.g., "Oh, I'm sorry, that was just a bad choice of language. The thing to keep in mind is ..." Rather than my having to say, for instance, "Yes, I believe we were wrong about that, but I couldn't convince my coauthors."

Another thing I guess I really didn't like—and this is only a new one for me, it's not something that had eaten away at me before—is the slogan "Gleason's theorem can be regarded as the greatest triumph of Bayesian reasoning." To say that is to imply that Bayesian reasoning IMPLIES Gleason's theorem. I don't think you mean that. I think what you mean is that it is one of the most valuable additions to Bayesian reasoning ever. But I won't cause trouble here: I know you both like the saying.

I think I read the paper very carefully again—more carefully than ever actually. I will live with its consequences, and I apologize for dragging my feet for three or four years.

Now, I just wish we three could come to agreement in our views of quantum mechanics!!!!!! So that I could look back on that day and say, "We few, we lucky few, we band of brothers ..."

20 September 2001, to Wootters, "Praise, Folly, Enthusiasm"

A million thanks for your encouraging letter of August 30! But also, please accept my apologies for not writing you back for more than 20 days. I had wanted to write you a rather long, contemplative note in response, and I kept waiting for the right mood. Somehow it just never came—first with my travels to Munich and Belfast getting in the way, and then with all the horrible events in the world in the last 10 days.

I certainly like aspects of this speculation of yours. Indeed, I wish you would write it down so we could all have a chance to think about it a more deeply. (Would you be willing to do that? You do have tenure now, and our individual lives are finite: we should never lose sight of that.) But of what little bit I understand of your ideas presently, they don't seem to have as much reciprocality or as much dynamism as I would like to think the world has. What I mean by this is that your ideas seem to carry a significant flavor of the Cartesian cut: the graphs take the place of res extensa, and the

identifications take the place of res cogitans. The two realms—as I understood your explanation—don't really interact: The graphs are timeless and independent of the identifications we might make between them.

But, as you say, we are both speculating. And we both realize that. Here's the way William James put it:

> The history of philosophy is to a great extant that of a certain clash of human temperaments. Undignified as such a treatment may seem to some of my colleagues, I shall have to take account of this clash and explain a good many of the divergencies of philosophies by it. Of whatever temperament a professional philosopher is, he tries, when philosophizing, to sink the fact of his temperament. Temperament is no conventionally recognized reason, so he urges impersonal reasons only for his conclusions. Yet his temperament really gives him a stronger bias than any of his more strictly objective premises. It loads the evidence for him one way or the other, making a more sentimental or more hard-hearted view of the universe, just as this fact or that principle would. He *trusts* his temperament. Wanting a universe that suits it, he believes in any representation of the universe that does suit it. He feels men of opposite temper to be out of key with the world's character, and in his heart considers them incompetent and 'not in it,' in the philosophic business, even though they may far excel him in dialectical ability.
>
> Yet in the forum he can make no claim, on the bare ground of his temperament, to superior discernment or authority. There arises thus a certain insincerity in our philosophic discussions: the potentest of all our premises is never mentioned. I am sure it would contribute to clearness if in these lectures we should break this rule and mention it, and I accordingly feel free to do so.

Since reading this a month ago, I've been wondering what my temperament really is. What is the "potentest of all my premises"? I haven't completely tied it down yet, but I think it has to do with the idea that the world can be moved (from within), that it is malleable, that it is still under construction. That the future, for better or for worse, is not yet determined. And that this malleability—like the turtles—goes all the way down: there is no ultimate level that grounds it.

This, if it is, my "potentest of premises" has so far taken manifestation in the speculation that the components of the world are "sensitive" or "irritable." But I suppose that is part of the "impersonal reasons" that James speaks of—i.e., that they form not quite the whole truth of my motivation. The better truth is perhaps that the components of the world—the things that come out of the ways we carve the world up—are movable in a deep sense. They reach out and affect us in ways that we cannot foresee, and we reach out and affect them in ways that they cannot. And through that intercourse, birth arises in the world in a sense every bit as real as biological birth.

But I don't know how to make any of that more precise. It stands at just a program and a direction, but one I take seriously.

Your letter gave me courage, and with such speculations, believe me, one needs courage! How I wish we could get together more often to hash these things out, letting our disparate speculations refine each other. I feel deep in my heart that there is progress to be made—technical progress—it's just a question of building a community with a critical mass of ideas, constructive opinions and techniques.

What are your plans for your sabbatical this year? In which direction do you plan to use your time? Will you be visiting IBM often? If I could get a chance to talk to you more often about the sublime side of physics, I'd surely take it.

By the way, let me draw your attention to a mistake I made in the paper you read. Howard Barnum brought it to my attention last week, and next week I plan to write a short comment on it

and post it on `quant-ph` (before someone else does). It is in the derivation of the tensor product rule. For the most part the derivation holds, but I got sloppy in the very last step. I.e., by requiring the existence of a noncontextual probability assignment to the outcomes of local measurements (with one-way classical communication), one does indeed get that the probabilities are controlled by a linear operator on the tensor product of the two spaces. But, these assumptions alone don't get you that the linear operator ought to be a positive semidefinite operator. That requires more assumptions. In principle, one should be able to state those assumptions as a restriction on the correlations one can obtain by local measurements, but I don't quite see how to do it yet. (An easy way out would be to require that the linear operator give rise to positive probabilities for all measurements, local and nonlocal. But, that's kind of a dull answer after relying on locality for so much. I'm sure one can find a more interesting answer.) For instance, I'm not quite sure how to tackle the question of ping-ponging measurements in this framework. Or even whether that's the sort of thing that should be looked at for a natural restriction.

There's so much work to be done! But, it's our place to do it ...

20 September 2001, to Mermin, "Pots, Kettles, and Frying Pans"

More seriously. I think the line you are talking about is not the line I was talking about.

Caves has been writing me things like this:

Cavesism 20: *I do want to conclude with my obligatory diatribe against wholly e-mail exchanges. You think all your messages were perfectly clear, I think all mine were perfectly clear, Ruediger probably thinks the same, but the evidence is that they were not. It's just not possible to come to agreement by e-mail, the reasons among others being that, first, questions arise in reading something after which further comments get devalued and, second, peripheral and main points often have their roles exchanged when a message is read.*

And ever more I am having to come over to his view on this subject. But, it is just so hard to give up my old email habits, especially since I've seen them become so ineffective of late. It's like a captain who just can't tear himself away from a sinking ship.

So let me say a few words in response to you. You should know I've got too much invested in the phrase "knowledge about the consequences of our interventions" to back out: it's part of my whole being. If you think that is what is going on with me, you are mistaken.

I do not deny: (1) that trails of evidence exist, and (2) that trails of evidence are created in part by our interventions into the world. What I deny is that those trails of evidence can ever uniquely determine a quantum state, even a pure one. That is what is at issue (with me at least), and that is what yesterday's note was referring to. (Somehow I get the impression that you saw something completely different in the note.)

A detector goes click. You write down that it went click; I write down that it went click. To that extent, the click is part of the objective world independent of us; it is part of the trail of evidence that you and Caves speak of. (Now, maybe the measurement apparatus itself isn't completely part of the agent-independent world—somebody built it to begin with—but that is a different story.) The formal structure of quantum mechanics says that we must identify the click with an element in some POVM. Fine. I agree with that. (Presumably you do too.)

The issue is, which POVM? And which state-change rule associated with that POVM? Show me a place in quantum mechanics where you are told how to do that. I say it is a subjective judgment, just like the quantum state is. Or, more precisely, it is exactly *that* subjectivity that keeps our view of the quantum state as being subjective from being an inconsistent notion. In any

given experiment, if there is a single POVM (and state-change rule) that is correct and objective, then so must be the quantum state.

Do you just not see the mapping problem here? If a POVM (and state-change rule) is an objective property of the interaction of two systems, then so will be the post-measurement quantum state ascribed to one of those systems. Now, if the quantum state is an objective state of affairs, why quibble about calling it "knowledge" or "information"? That would be just using two words for it that never needed to be invoked in the first place. If the quantum state is objective, then call it "quantum state" and be done with it.

If you can stomach it, try to read my note "Practical Art" again. It was meant to be a call to do some soul-searching. It was meant to try to persuade you that the subjective element can never be eliminated in a theory that makes fundamental use of probabilities. It was not meant to convey the idea that one cannot draw a line between systems and apparatuses. It was not meant to be a call to join the ranks of the Everettista.

Chris (on a rainy day—it always happens on a rainy day)

20 September 2001, to Caves, "Re: Comment on Practical Art"

You still don't recognize that the difficulty is a logical one, do you? There is nothing squishy, postmodern, deconstructionist, or new age about the issue: if POVMs (von Neumann ones in particular) are nonsubjective, then so are quantum states. Period. You can't get around that. If you want to claim that quantum states are subjective judgments, then you have to accept that POVMs (and their associated state-change rules) are subjective judgments too. Else the post-measurement states that they give rise to would be objective.

I say this slightly better in my reply to Mermin. It is not that I am leaving the realm of science: it is the strict practice of the art that put me here. Assumptions $\longrightarrow$ conclusions.

Don't worry though: I'm not offended by your note. I'm just (continuingly) surprised by your immense rigidity.

25 September 2001, to Caves, "Further Comment on Practical Art"

Cavesism 21: *I was initially mightily annoyed by the tone and content of your message. My comment said nothing about and intended nothing about being "squishy, postmodern, deconstructionist, or new age." Moreover, your reply didn't address any of the points I raised, apparently because they were all just further examples of "rigidity."*

Let me reply to your last point first. Here's what happened: I thought I'd reply with a short and to-the-point comment first (about what I saw as the overriding issue), and then follow that up with a detailed reply to your particular points. But then time ran out before having to leave for a weekend trip to the beach. I certainly did not mean to ignore your points, and I certainly won't (as time permits this week). It was just that I was banned from email for the weekend: I'm sorry I didn't warn you of that.

Now, as far as the postmodern business, I apologize for being oversensitive on that point. It had to do with this remark:

Cavesism 22: *It suggests to me that agonizing over the sorts of questions in your hag-beauty-insurance story is not going to help us understand what we're doing as scientists and, moreover, that it will lead our program into a place where no other scientists wants to go.*

This is (what I recall to be) the third time you have made such an allusion, i.e., that I am going places where no scientist ought to want to go. The only conclusion I could draw from this was that you were implying that the subject of this whole discussion is nonscientific. I hope you will at least understand that I might find that insulting. And my reply was an attempt to put a quick end to that train of thought. Over and over, I feel that I have been trying to make a simple point, and most importantly, a logical point—it just happened to be an unpalatable point. So, in the end I guess it seemed to boil down to both of us accusing the other of being unscientific.

I will write a (calm) set of emails replying to all the new issues you and Ruediger have brought up in the next day or two.

03 October 2001, to Mermin, "Kid Sheleen"

What a busy couple of weeks it's been for me! I am sorry it has taken me so long to reply to you. But cars needed buying, lots of newspapers needed reading, lots of soul-searching needed doing (before canceling what would have been my fifth trip abroad this year), and Bell Labs needed some tolerance (with the transition from Brinkman to Jaffe as our Research VP).

Merminition 28 : *Sorry. Didn't mean to go on for so long. Caves is right; this is a rotten way to have a conversation. No need to reply.*

Yeah, Caves is right in many ways. I've never had more frustration in an email run before. But then, I don't think I've ever made such a headlong transition in open view before either. Regardless, however, I have certainly benefited from this: I've never had to strive so hard to make a simple point clear, and I think that gave me a load more perspective on the issue than I would have come to by myself. Strangely, in a way, the whole affair has hardened me and made me confident in this new direction of thought.

Merminition 29 : *so I'm sympathetic to what you're saying, but I worry that you're giving up "objectivity" on too many fronts.*

You think you worry! That reminds me of a dialogue in the old movie *Cat Ballou* (with Jane Fonda and Lee Marvin):

> Jackson Two-Bears: Kid, Kid, what a time to fall off the wagon. Look at your eyes.
> Kid Sheleen: What's wrong with my eyes?
> Jackson Two-Bears: Well they're red, bloodshot.
> Kid Sheleen: You ought to see 'em from my side.

But I don't think I've given up objectivity on too many fronts: If there's one thing for sure, it's that I don't want to go too far. The tenet I hold fast to is that there is something happening in the world, something was happening before scientists appeared on the scene, and something will continue to happen (in one form or another) should we wipe ourselves off the planet next year.

The issue is, how does that something interface with us, to what extent can we grasp it, and how do we modify it by our very attempt to grasp it?

Let me try to reply to some of the points in your letter.

Merminition 30 : *On the other hand you did seem to be undermining the impact of the trails of evidence when you gave a central role to the question of whether your Lucent colleague really had the polaroid oriented the way you had been led to believe it was oriented.*

Hideo Mabuchi is a young professor at Caltech (probably the youngest); he won a MacArthur Fellow (i.e., $500K) last year too. He started grad school at Caltech at about the same time that I started up with Caves. You can see he's the smart one. We've been friends ever since we first met at the Torino meeting of 1994. He keeps the napkin where I first explained the Holevo bound to him; I mention him whenever I can (like in the NATO paper) because of all the pearls of wisdom he gave me.

Merminition 31 : *I'll grant you that that's something to worry about, but it seems to me on a different level from the characteristic quantum ambiguities,*

Three months ago, I would have thought it was on a different level too. But now it is clear to me that the two NECESSARILY feed in to each other. (Though the whole issue has been building in me for two years: See Samizdat, page 127, in a note to Caves titled "The Dangers of Probabilismo." I'll talk about that more in a minute.)

Merminition 32 : *unless you want to deal with your colleague and his polaroid on the same footing as the photons, whence my joke about your closet Everretism. (Please read on before concluding that I still haven't got the point.)*

I did read on—many times—and I think you've got some of the point, but not the whole thing. In particular, I still contend that Everettism is the antithesis to my point of view.

There is a sense in which Hideo and his polaroid are on the same footing as the photons (and I have always thought this). It is that in reasoning about them—i.e., in reasoning about what kinds of traces they will leave from my interactions with them—I am obliged to use the formal structure of quantum mechanics if I want to do the best job I can in that reasoning. But I don't think your remark in particular has anything to do with this.

Your remark seems to be more of the flavor: Chris says POVMs (even simple von Neumann measurements) are "subjective judgments," so he must mean that their outcomes are every bit as dreamlike and subjective too. Measurements that have no concrete outcomes? What else could this be but Everettism in disguise?

But there is a non sequitur there (if indeed that is the flavor of your reasoning). It is that you (or my caricature of you) think I am identifying—in a supremely steadfast way—the "click" that makes its way to our senses when we perform a measurement and the very being of a POVM.

Instead, the idea is that the "click" is REAL, as real as you could want anything to be (for those concerned, for those who know of it). But, the index b (from some POVM $\{E_b\}$ that we associate with it) is not the thing-in-itself. It just gives the "click" a NAME and a CONTEXT through which we can draw inferences and through which we can start to contemplate further reasoning. It is that identification which is an ultimately subjective element; it is not the "click" itself.

I had hoped to draw your attention to this distinction with the Penrose beauty-hag example, but I see I completely failed on that count (not only with you, but also with Caves and Schack). There is the stuff of the world—the "click" (in part). And then there is our description of it. The two things are not the same things; one lives in my head, and one lives partially outside it. To accept quantum mechanics, is to accept a template for that description and to accept a method of manipulating one's judgments thereafter.

Does any of this make sense? I think it is the main point I have been wanting to make to you, but as far as I can tell, you've ignored it.

Merminition 33 : *I agree that figuring out what measurements are associated with what operators (even at the von Neumann level) is something you have to bring in from the outside. So is knowing*

what the Hamiltonian is. I'm not so sure these are the same kinds of subjective judgments as a state assignment, as you maintain next. . . .

The new thing I've been saying since August is that the identification of a particular POVM (and one element therein) with a measurement "click" is a subjective judgment. I had not clearly appreciated that before. However your mention of the Hamiltonian is apropos, because that one, at least, I had caught before. (That is what Samizdat, p. 127 is about.) There are several ways to see that "the Hamiltonian" must be on the same subjective footing as the quantum state. Here's one; I'll just quote you:

Merminition 34: *I worry that "objective" is taking on too many different meanings here. For example, EPR makes it clear (at least to sensible people) that the polarization state of a photon is not an objective property of that photon.*

Will you accept that the existence of quantum teleportation carries as much force as the EPR argument that the quantum state is not an objective property of a photon? If you will, then can you tell me what the import of all the recent papers on "teleporting a unitary operator" is? (See for example `quant-ph/0005061`, but there are a plethora more.) But even if you don't, the argument is simple: Hamiltonians can be toggled from afar by our measurements on entangled states. (Carl gets himself out of this by saying that the only thing being toggled is the "effective" Hamiltonian, not the underlying quantum circuit, but I say where there is a tear in the fabric there is a rip.)

Merminition 35: *I think you're saying that if the vertical alignment of the polaroid is an objective fact, then the state — "vertically polarized" — of the emerging photon is also an objective fact. But that's not the same as saying it's an objective property of that photon.*

That is right, there is a difference. And you have gathered what I was saying (almost) correctly. (I wrote "almost" to help remind you of the points above about identifying "clicks" with POVM elements.) The point is, the quantum state had better not be an objective fact, or the point of view that Caves and Schack and I have been trying to build up will be in deep trouble. What is wrong with taking a quantum state to be an objective fact, as long as one drops the insistence that that fact be localized with the photon? (I.e., as long as one does not make it a property of the photon where it stands.) At first sight, maybe nothing: I think that is probably the point of view you are trying to build; it is also the point of view Philippe Grangier has been trying to build in his recent `quant-ph`'s. But, on second sight, one cannot forget that the quantum state uniquely specifies a set of probabilities. If the quantum state is an objective fact, then so are those probabilities. And now it is on your shoulders to tell me what objective probability can possibly be. I won't stand for anything short of an operational definition.

Merminition 36: *I worry that "objective" is taking on too many different meanings here. For example, EPR makes it clear (at least to sensible people) that the polarization state of a photon is not an objective property of that photon. It appears from the above that you believe you can only consistently take this position if you deny objective status to the outcome of the actual polarization measurement which enables you to predict the outcome of the measurement that has not yet been made.*

No. See the point above, where I used a little TeX notation. I accede to the objective status of something happening in a measurement intervention. I just don't accede to an objective status for what we decide to call it, i.e., for which POVM we decide to associate it with.

Merminition 37: *(Forgive me, but this smells like a many-worlds strategy again. Your answer to the EPR paradox would seem to be to deny that the first measurement had an objective outcome. Recall Henry Stapp who has been saying for decades that only for the Everretista is nonlocality not a consequence of EPR.)*

I forgive you. But I hope you'll tell me that, with your newfound enlightenment, it doesn't smell so much like Everett anymore. (Your imagery conjures up my own imagery of walking on a warm day near a trash can full of lobster parts behind some coastal New England restaurant.)

Merminition 38: *I'd prefer a middle ground which allows me to talk about objective facts but not objective properties. (It just now strikes me that one might call this correlations without correlata.) I've never been sure I can do this coherently, so I'm sympathetic to what you're saying, but I worry that you're giving up "objectivity" on too many fronts.*

I am in partial agreement with your first sentence, and I would like to think I have hit a sweet spot for that part. I'll give you objective "clicks" (though I might not call them "facts" ... but that's a story I probably shouldn't get into right now); I just stand fast against the objectivity of the quantum state. What is more middle ground than that? However, I do not share your aversion for objective properties.

It seems to me, quantum systems do have some properties that we can get our hands on. I usually preach the bundle of information-disturbance curves associated with a system, but let me try from a different angle to convince you of at least one property. I say that the quantum state cannot be an objective property because we can toggle it from a distance. I say that the Hamiltonian cannot be an objective property because we can toggle it from a distance. But what about a quantum system's Hilbert-space dimension d? Can you think of a way to toggle that number from a distance? I can't. And so, to that extent, I'm willing to call the raw number d an objective property of a part of the world. Now, what is the physical meaning of d? Well, that's why I struggle with all this information-disturbance stuff, but that's not the issue at hand. The issue is that one need not give up on all objective properties.

There are things in the world beyond our control: One them is the outcome of a quantum measurement, and one of them may be the dimensionality of a Hilbert space. Objectivity means nothing to me if it doesn't mean that some things are beyond my control, are beyond my whim and fancy. To the extent that I'm willing to say this, I don't think I'm giving up on objectivity on too many fronts.

Does this strike any chords in you?

PS. By now you should have received our modified version of "Making Good Sense." I won some good battles there: We no longer claim that two observers must be compelled to the same unique state via a Dutch-book argument. But I lost some too. I continue to think the paper is misleading as hell, always talking about a "unique" state assignment and using the word "certainty" in a way that still troubles me. We were able to compromise only in that I thought things were now worded in a sufficiently vague way that I could worm out of them in my future talks and publications. I don't think we say anything factually against my beliefs, but the reader will have to be on his toes to not get fooled about where I really stand.

PPS. Here's another thing I ought to tell you. PRA made the mistake of asking me to referee the BFM paper. Despite what I wrote you earlier about probably accepting the challenge if it came up, I decided to decline the opportunity. I like the math of the paper, but I just could not agree with what you make of it. It seemed more appropriate to let some less tainted souls than mine tell you what they think of it.

03 October 2001, to Caves-Schack, "Replies on Practical Art"

Cavesism 23: *The philosophers tend to proceed by telling a story—reasoning by analogy, they call it. The actual problem is too hard for them to formulate, so they immediately introduce a simple analogy, reach a conclusion they like within the analogy, and then transfer the conclusion back to the actual problem, without ever justifying why the analogy has anything to do with the actual problem.*

Except for omitting the final justification—which is more than important—is this so different from what you teach in your physics classes? I.e., That one ought to try one's ideas out on a simple example first? One that may already contain the essence of the problem, before embellishing it with too many details?

Cavesism 24: *I enjoy reading your stories, but perhaps you're falling into the same sort of trap in a less obvious way. The difficult and very personalistic questions about assigning probabilities in your hag-beauty-insurance story are important in thinking about Bayesian probabilities. These personalistic factors are well known to be present in a subjective interpretation of probabilities, but do we really have to worry so much about them in the context of interpreting quantum probabilities?*

My point was to remind you guys and Mermin that these personalistic factors always must exist, else we would have no need to take such pains to talk about a "subjective interpretation of probabilities." If they are well known (as you say), then they should not be forgotten and replaced with purely objective "trails of evidence."

The point is, yes we must worry about them in the context of interpreting quantum probabilities. We must recognize that that is part of the very problem. Once we have recognized it, then we can move on and almost forget that the issue was ever there—just as one can do in whole textbooks on orthodox probability theory—but that first step is a supremely important step.

Cavesism 25: *You and others write papers every day where this party assigns this state and that party assigns that state, and I don't see any of these papers agonizing over difficult, personalistic questions of what state to assign. You're right to keep badgering us to pinpoint why this is true, but the fact that it is true—we don't worry about this kind of stuff when making quantum state assignments—leads me to believe that there is an answer.*

Nor do you see any sophomore-level textbooks on probability theory agonize over these personalistic questions when posing its exercises at the end of each chapter. On the one hand, you completely missed the point I was trying to make, but on the other you also completely answered it.

The point is one does not have to worry about these personalistic questions to get quantum mechanics as a *calculational* tool off the ground. In that regard, the present issue is no different than with classical probability theory. Indeed, this is probably why in both theories a large sect of the practitioners have turned to "objective" interpretations of their main terms (alternatively, probabilities and quantum states) in such a misguided way. It is in the very recognition that personalistic questions exist, that one is compelled to finally get the foundations of the two theories straight.

In practice, what almost always happens? In the case of classical probabilities, when given a specific problem, one reduces and reduces the problem until one has transformed it into an equivalent problem for which one feels confident in making the uniform probability assignment. Thereafter one *derives* the probabilities for the problem of real interest by transforming and grouping, etc., until one rearrives back at the starting point. (This is a point I probably first learned from Ruediger.)

Think for instance, if I were to ask you what is the probability of obtaining a 7 or an 11 in a roll of two dice. Your mind would probably first jump to the judgment that all six sides of each die are equally likely, and then let the calculations flow.

Now of course, being a Bayesian, you would leave open the possibility for something else than a uniform assignment in that step above. But in practice, there are some things that most of us can usually agree upon ... and those are usually the starting points for textbook problems.

The issue is little different in quantum mechanics. When presented with a problem of calculating a quantum state for a given physical system, what do we usually do but reduce and reduce (or expand and expand) the problem until we come across an equivalent one for which we are confident we can predict the outcome of some measurement with certainty? Thereafter we work our way back just as before. Just think of Scully and Lamb's derivation of "the" quantum state of a laser. Alternatively, think about Moelmer's justification of the same state. [Indeed one might say that this is what the whole (worthless) decoherence program amounts to: deriving one subjective state from another and then thinking there is something deep about it. But that's an aside.]

Now just as before, being a Bayesian, one ought to leave open the possibility for something else than the particular pure state in the basic step of this derivation. But in practice, there are some things that most of us can usually agree upon ... and these are usually the starting points for textbook problems.

I think the similarity is overpowering. It is enough in both cases to recognize that ultimate personalistic issues exist, but then the homework assignment goes on—the student reduces the problem to a judgment few people would dissent on.

Cavesism 26: *It suggests to me that agonizing over the sorts of questions in your hag-beauty-insurance story is not going to help us understand what we're doing as scientists and, moreover, that it will lead our program into a place where no other scientists wants to go.*

Looking back over this note again, your language really was very scolding throughout—"trap," "badgering," "agonizing," "a place where no other scientists wants to go"—in spite of the fact that you warned me it would be "a constructive and gentle criticism." I understand that I am guilty of no less: There is no doubt that I can be arrogant and abrupt (and paranoid) at times. But in all this massive email, I feel that I have been providing a service, sharing ideas that I might not have if I didn't feel we should be brothers in arms. It became a little hard to gulp that all these notes might be viewed as little more than a nuisance.

Cavesism 27: *The answer might be as simple as this: we can only do science in situations where we scientists have agreed that such personalistic factors can be essentially eliminated, and quantum mechanics is the very pinnacle of this kind of situation. I think that's the content of our statement that "Gleason's theorem is the greatest triumph of Bayesian reasoning" and of our "principle of quantum indeterminism."*

I think the answer might just be as simple as that, but at a level higher than the one you are contemplating. The agreement we need is in accepting quantum mechanics as a method and a restriction for shuffling about our more mundane, everyday beliefs. What that entails is accepting POVMs as the structure of the questions we can ask a system and the Kraus state-change rule as our method for updating our beliefs.

Everyone keeps asking, what is the objective piece of quantum mechanics? I answered some of my beliefs on that issue in the letter I just sent off to Mermin (and then forwarded to you). But I think there is also quite a bit to be learned on the issue by first turning the question toward

Bayesian probability theory. What is the objective piece of Bayesian probability theory? I think all three of us are in agreement that it is not in the particular probability assignments that one might make. But is there *any* objective piece at all? I think there is. Take Bayes' rule as an example. I would say that it is something objective in the theory: it is the ideal of behavior. If one doesn't use it, one can be taken advantage of. You agree that Bayes' rule is the ideal of behavior, and I agree that Bayes' rule is the ideal of behavior: it would remain the ideal of behavior if all of us were wiped off the planet.

Likewise, it seems to me, Gleason's theorem plays a similar role. There must be a sense in which accepting that the structure of our questions to the world (or, alternatively, our interventions upon it) conforms to the structure of the POVMs must be the ideal of behavior—something not so far removed from Bayes' rule itself. It is the ideal of behavior in the light of some crisp, physical fact. I don't know what that fact is yet (in any precise sense), but that does not stop me from seeing the outline of how the various structures in quantum theory should be classified:

measurements = POVMs	objective feature (physical fact)
Born RULE (via Gleason)	objective feature (an ideal of behavior)
Kraus state-change RULE	objective feature (an ideal of behavior)
quantum state	subjective judgment (always)
time-evolution map	subjective judgment (always)
Hilbert-space dimension	objective feature (physical fact)
particular POVM assignment	subjective judgment (always)
particular Krausian assignment	subjective judgment (always)

The list is not exhaustive; but I think these are the ones I see clearly at the moment.

The point is: Agreement required for science? Yes. Compelling interpersonal agreement as a (potential) statement about the agent-independent world? Yes. Agreement necessary at the level of quantum states? No.

Cavesism 28: *I'm not sending this to Ruediger and Mermin, but you can send it to them if you think it's worthwhile to do so.*

Well, clearly I thought it was worthwhile to share my answers ... but there's that arrogance again. ;-)

With a smile and a conciliatory tone,

03 October 2001, to Caves-Schack, "Further Replies on Practical Art"

Cavesism 29: *I think we agree that there are things that are effectively facts in the effective reality of ordinary experience.*

In the words of Bennett's father (in such a context), "These are very deep waters." Since becoming enamored with James, Dewey, and Schiller—and having read copious (by my standards) amounts of them—I'm not completely sure how I should answer you. The issue is, I'm not completely sure in which sense you are using the word "fact." I have a feeling it is a more loaded sense

than you would guess. But I don't want to get into that now: You suggested some simmer time, so I will leave you some until it becomes absolutely necessary.

There is, however, one thing I dislike about this sentence, and that is the phrase "the effective reality of ordinary experience." But you touched upon that very point in your note "More on Pots and Kettles"; so I'll say more to the issue in detail when I reply to that note. The main thing, though, is that I would say our ordinary experience is the rawest stuff around: It's the very stuff from which we build these super-smooth pictures by way of which we derive our further expectations. There is nothing effective about it: It is the stuff, it is the starting point. To use the word "effective" makes it feel secondary and derived (which is what you have been striving to get at, not me).

Cavesism 30: *The questions arise in what those facts tell us or, perhaps, in whether and what they compel us to believe. The argument is about pure-state assignments, not about mixed-state assignments. You believe that the subjectivity of pure states requires that it be possible for different agents to assign different pure states.*

Yes.

Cavesism 31: *To say something is subjective is to say that it exists only because of us and does not have an independent existence out there in the world. It also implies that different agents can disagree, the degree of possible disagreement being just the flip side of the degree of intersubjective agreement.*

Yes.

Cavesism 32: *Suppose we had the idea that facts in the effective reality force one to a particular pure-state assignment. The resulting pure state is then based on a trail of evidence in the effective reality and is embedded in each agent's mind. Is the pure state then out there in the world, independent of us?*

I would say, yes it is. The agent's state of belief is then an unneeded complication in everything under discussion. The fact is that there is a one-to-one correspondence between (sets of) facts and quantum states. You can say the agent's mind is nevertheless needed to "house" that state, but then, to me that looks to be nothing deeper than invoking the luminiferous ether to support the electromagnetic field.

Cavesism 33: *Do I have it right that this is the issue, or at least an issue?*

Yes. I have always perceived this to be the main issue. To the extent that I have said words all around this, it has been—I believe—to attempt to give different angles for viewing the same thing. Every time I saw you, Ruediger, Mermin, Brun, etc., be reluctant to accept the point, I tried to present it from a different angle so as to be more convincing. I take it now that everyone only found that to be confusing. But what else could I do? And I can't complain too much, because I think the whole process has sharpened my presentation of the point (which I maybe only dimly perceived at the beginning).

Cavesism 34: *I don't necessarily see where the pure state is if it's thought to be out there in the world. The trail of evidence is not a pure state; we construct the pure state from the trail of evidence, but dogs don't and dinosaurs didn't.*

A one-to-one mapping is a one-to-one mapping. I do not see how YOU cannot see that making these statements is not a tacit acceptance that the quantum state is an objective entity after all. Maybe you have thought this all along. Namely, that when you said a quantum state is not a state of nature, what you really meant was that it was simply not localized on the physical system it is meant to describe. It is a state of nature, i.e., it is a collection of facts within nature, it is just not living on the system it's intended for. From this point of view, it's clear why dog's don't use them: Dog's aren't clever enough or advanced enough technologically to discover the true states of nature.

But I surely never thought this when I used the slogan, "a quantum state is a state of knowledge, not a state of nature." If facts can uniquely determine a quantum state, and facts live in nature, then a quantum state is a state of nature after all.

Cavesism 35: *The pure state is not out there in the system for the reasons we have long discussed: the system can't report its pure state, and a system's state can be changed to any pure state drawn from incompatible sets without ever getting close to the system. It looks to me like the pure state is purely in our minds.*

I'm not sure how this remark fits in. You might be making a call for me to consider putting the "trails of evidence" into the mind, but I'm not sure.

Cavesism 36: *You are insisting, I believe, that in order for a pure state to be subjective, it must be possible for different agents to disagree on a pure-state assignment. You say, I believe, that if we are forced to a particular pure-state assignment by the facts in the effective reality, then the pure state becomes objective.*

Yes.

Cavesism 37: *I don't know where to come down on this. It is one aspect of the question I always ask of which aspects of maximal belief get translated from realism to quantum mechanics. It also has to do with the nature of the "facts" in the effective reality and thus how the effective reality arises out of quantum mechanics (this is the content of Mermin's Pots and Kettles). My own take at present is that the effective reality is a form of intersubjective agreement.*

Fair enough that you don't know where to come down on this: I will try not to lose my patience any more.

Cavesism 38: *You are right in principle that nothing compels us to a particular pure-state assignment, but clearly wrong in practice. All our experience with quantum mechanics suggests that we have no trouble agreeing about pure-state assignments, so no matter how the facts in the effective reality arise, there is nearly total intersubjective agreement on what they imply for pure-state assignments (this is the content of my Comment on Practical Art and further comments on Pots and Kettles).*

I hope my previous note addressed this adequately. In contrast to what you say, I believe that I am right in principle and right in practice. You might have said the "all our experience" sentence about classical probability theory if your name were Richard von Mises. He would have said that all our experience with dice shows that we have no trouble agreeing that its outcomes are all equally probable. But you're not von Mises, and you've had the advantage of having had 75 years of good Bayesians clear the air for you. The issue you bring up has little to do with the structure of the physical world, and little to do with the structure of the Hamiltonians we feel compelled to describe it with.

Cavesism 39: *Let me know if I have got your position straight.*

I think you did.

Cavesism 40: *If I have, then it seems to me that we are not far apart, the only gap being how much we are willing to ascribe to the apparent agreement that exists in assigning pure states. You prefer to emphasize that nobody can be coerced into this agreement, and I prefer to emphasize that in practice nobody has to be coerced into it.*

A point of emphasis can make a huge difference in a philosophy. And a difference in a philosophy can make a huge difference in the practical and applied questions one might ask of quantum mechanics.

And I'm off to Lupé's for the best Mexican food in New Jersey. (You know that's not saying much.)

More tomorrow!

04 October 2001, to Caves, "Finicky Sins"

By the way, in saying this yesterday,

> There is, however, one thing I dislike about this sentence, and that is the phrase "the effective reality of ordinary experience." But you touched upon that very point in your note "More on Pots and Kettles"; so I'll say more to the issue in detail when I reply to that note. The main thing, though, is that I would say our ordinary experience is the rawest stuff around: It's the very stuff from which we build these super-smooth pictures by way of which we derive our further expectations. There is nothing effective about it: It is the stuff, it is the starting point. To use the word "effective" makes it feel secondary and derived (which is what you have been striving to get at, not me).

it dawned on me afterward (on my drive home) that I was committing a sin: There was a time when I liked the phrase. As you know, I commandeered it when writing my *Physics Today* articles with Asher.

So, maybe I was being overharsh, or maybe just finicky. In any case, even when I used it unreservedly, I know that I had a distinct flavor of the phrase in mind from the way you had been using is. [See our discussion starting on page 133 of the Samizdat.]

Maybe I'll say more about this later today.

04 October 2001, to Caves-Schack, "Replies on Pots and Kettles"

This letter is going to be a hard one to reply to, because I don't quite see how much of what I have said in the past led to the points you make here. So, let me just plunge into the thing and see what comes out.

Cavesism 41: *I thought you were in the camp that holds that the our experiences—our actions or interventions and our perceptions of the world's response—are primary*

I thought I was too.

Cavesism 42: *and that the function of science is to account for them.*

It's this part of the sentence that I'm not so sure of (though it's not clear to me exactly what you have in mind). If our experiences are primary, then it does not seem to me to be within science's purview to account for them. I believe, instead, the view I have had for quantum mechanics for some time is best mimicked by these words I picked up from William James last month:

> Metaphysics has usually followed a very primitive kind of quest. You know how men have always hankered after unlawful magic, and you know what a great part in magic *words* have always played. If you have his name, or the formula of incantation that binds him, you can control the spirit, genie, afrite, or whatever the power may be. Solomon knew the names of all the spirits, and having their names, he held them subject to his will. So the universe has always appeared to the natural mind as a kind of enigma, of which the key must be sought in the shape of some illuminating or power-bringing word or name. That word names the universe's *principle*, and to possess it is after a fashion to possess the universe itself. 'God,' 'Matter,' 'Reason,' 'the Absolute,' 'Energy,' are so many solving names. You can rest when you have them. You are at the end of your metaphysical quest.
>
> But if you follow the pragmatic method, you cannot look on any such word as closing your quest. You must bring out of each word its practical cash-value, set it at work within the stream of your experience. It appears less as a solution, then, than as a program for more work, and more particularly as an indication of the ways in which existing realities may be *changed*.
>
> *Theories thus become instruments, not answers to enigmas, in which we can rest.* We don't lie back upon them, we move forward, and, on occasion, make nature over again by their aid.

Science does not account for our experiences. Science builds on them and gives us a structure by which to imagine pushing them to a new extreme. This is why I have laid such emphasis on calling the world "malleable" (for instance in my "Activating Observer" document that I shared with Ruediger ... and maybe you too, I can't remember). It seems to me, science does not say so much about what is, but what can be (subject to the limitations to our actions captured in the very structure of the given theory).

Cavesism 43: *Trouble is that when our interventions proceed to too fine a level, the world's response is not deterministic and, furthermore, cannot be described within the realistic language of ordinary experience. Surprisingly we find that we can use the strange, unrealistic formalism of quantum mechanics to describe the intrinsic indeterminism that intervenes between our actions and our perception of the world's response.*

What do you mean by "too fine a level?"

Cavesism 44: *I thought you were ascribing some sort of objective or intersubjective reality to our primary experiences.*

Pretty much. Or, at least that's what I thought too.

Cavesism 45: *I thought the difference between you and me was that I think that we must somehow derive from quantum mechanics—or, at least, make consistent with quantum mechanics—the apparently realistic features of the emergent "effective reality" of ordinary experience, whereas you think this is unimportant, thus accounting for our different reactions to the decoherence program. But you now seem to be demanding much more than I do,*

I've always thought that I've demanded much less than you, and I don't think I've changed my tune on this account for several years. For the view I dream of constructing, what is the classical world? It is a world for which the agents describing it are full of ignorance and the best to which they can muster is a lot of imprecise control.

I look out at one of the trees outside my window, and I ask how I might capture everything I'm willing to say about it into a single density operator. I can't say much about that, but I'm willing to bet that if I would carry the project through, what I would end up with will be so mixed, so thermal, that it'll be just about commuting with anything else I might have come up with, even if I had stared at the tree a little longer. This little fact—it seems to me—has nothing to do with the Hamiltonians of the world (as if they were objective things). It is a function of my pure ignorance and my unwillingness to tear the tree apart and refine my beliefs.

The idea toyed with here is that it is just ignorance, no matter how we each walk into the room with it, that leads to the classical world. If I am so ignorant as to use an almost commuting set of density operators for a given object, then any (gentle) attempt you—as another scientist—may make to refine those beliefs will not cause my beliefs to be any less valid: Your information gathering, will not cause a disturbance to my description. And therein—the speculation goes—lies the essence of classical mechanics.

Cavesism 46: *that before we even start, we be able to explain exactly how the effective reality works and at what point it emerges. As David points out, this is exactly what the Everettistas demand.*

As best I can tell, this remark can only come from viewing my program (more accurately, my dream) from your philosophical predispositions. I would never demand that we "explain exactly how the effective reality works and at what point it emerges." The classical world comes first. Quantum mechanics (as a theory of inference) extends beyond it, by taking into account new phenomena that simply can only be seen when working in a regime of less ignorance.

A relevant ditty to read might be my essay "Always One Theory Behind" on page 464 of the Samizdat.

Cavesism 47: *Of course, after making the demand and finding present responses unsatisfactory, you and the Everettistas go in quite different directions. They, out of an anal need for naive realism, simply concoct a naive realism to go with the state vector. You certainly aren't going in that direction, but being risk averse and already burned, I'm not going to risk a description of your direction here.*

I loved the phrase "anal need."

Cavesism 48: *To your credit I think you won't claim to have gotten your ideas worked out entirely (perhaps the rest of us can be allowed some access to that defense). Still you might want to think about the road you're traveling on and how it relates to this question of taking as given the apparently objective experiences of our daily lives.*

Thanks for the credit. It's refreshing to be complemented for not acting like a guru with all the answers. People at foundations meetings will have none of that. It's been my experience that they demand you tell them exactly what reality *is* ... before they quickly tell you you're wrong. (Matthew Donald told me he couldn't take me seriously as a foundations researcher because I keep evading the question of what reality *is*.)

Anyway, in conclusion, give me some feedback: Did I answer anything that you wanted answered? (I sure hope I did: I'm trying.)

04 October 2001, to Caves-Schack, "Replies to Morning Coffee"

Schackcosm 36 : *Oh no, I thought we had reached some agreement. One problem is that Kraus operations (unless they are 1D projectors) will not bring you to a unique place. Which means that qm does NOT provide a universal rule of coming to agreement. You need some ad hoc assignment.*

Yes, I had always understood that. I am sorry if I was sloppy about expressing it, but I thought I had always emphasized that there are two things that one can contemplate: 1) if Alice and Bob have complete freedom to choose what measurement they might perform, and 2) if instead they have at their disposal some fixed measurement (perhaps not of their choosing). In the later case, only certain initial states for Alice and Bob will lead to further agreement after the measurement. (See, for instance, my note to Mermin and company dated August 7.)

Schackcosm 37 : *The other problem is that your conversation is far too playful. State assignments are compilations of betting odds. They are COMMITMENTS. Chris in your dialogue should have been deeply shaken. He would have betted his house in New Jersey on this outcome to be impossible.*

Yes, perhaps. But, on the other hand, there is a counter trend in you that troubles me. And that is the basic philosophy that comes across as the message, "Once a quack, always a quack." What I mean by this is, suppose I ascribe a pure state $|\phi\rangle$ to some system, whereas you ascribe $|\psi\rangle$. As we have laboriously teased out of this correspondence, from your perspective, I am simply wrong. My judgment is not to be trusted (from your perspective). This much we agree upon. But I sense that you want to conclude more: Not only am I not to be trusted in my conclusions about the given system, but that I am not to be trusted about anything. You conclude that I am truly insane just because I adamantly disagree with you about one thing (as captured by our differing pure states). I say that goes too far.

I tell too many stories, but here is a true story. In discussing the cardinality of the natural numbers versus the even numbers, Kiki will accede that there is a one-to-one and onto mapping between the two sets. Nevertheless she contends that there are more natural numbers than even numbers. I have never had more annoying conversations than the one we revisit about once a year on this subject. I simply cannot convince her that she is not being logical on this issue. But still I do find that I trust her judgments on other issues.

A "misstep" on a quantum state (even a pure state), it seems to me, is not the end of the world precisely because of this.

Granted, I am a bit confused on what I think the ascription of a pure state actually does capture, but I think making it carry the weight of an agent's rationality or irrationality goes too far.

Schackcosm 38 : *Yet another comment: You have said nothing in all your notes (to my best knowledge. . .) that tells me why this situation is different from Chris being certain that there are two chairs in this room, and Carl with Hideo's help convinces him that he was wrong. I'd say either Chris was tricked, or he had hallucinations. Todd said: "This is why we say that insane people 'out of touch with reality' ". I said the same thing in a different way.*

Since you ask this question more pointedly in another note, I'll wait on answering it until I get there.

Schackcosm 39 :

> CAF Said: So, in saying this, I don't think one has to do anything fancy with a measurement to get this thing you call fuchsian consistency off the ground. No intricate constructions of states seem to be needed. Choose ANY (mutually agreed upon) von Neumann measurement you like, and Carl and Chris will have to agree at the end of the day. One of them will feel foolish for having been so wrong, but as long as he is rational, agreement can still be had. As Ruediger said, the idea is that this is a property of quantum mechanics (in its capacity as a law of thought), not of the initial states.

I found the idea quite attractive that in a quantum world, differences can be resolved (using a well-chosen measurement) that would be impossible to resolve classically. I thought that this was what you had in mind.

As far as I can tell, no physical statement (no ascription of a phase space point) is impossible to resolve classically. What is different is that quantum mechanics can do that even without the preexistence of a phase-space point ... and that surely is a property of quantum theory. But I said I'd come back to this in another note.

04 October 2001, to Caves-Schack, "Replies to a Conglomeration"

Now let me reply to a conglomeration of notes from you two.

First to Ruediger:

Schackcosm 40 : *I remember you writing something to the effect that the click in a measurement is the closest thing to a fact one could come up with (sorry for not looking it up, but you write TOO MUCH).*

Did the reply I wrote to Mermin yesterday make any sense to you? I am now in the habit of drawing a distinction between a "fact" and a "proposition." The difference was not so important classically, but I now think it is paramount quantum mechanically. The fact (or consequence of our intervention) is the raw, uninterpreted, unclassified stuff of the world. It is the real stuff that makes its way to our senses and then to our brain to be pondered. The proposition, on the other hand, by its very nature attaches a meaning to the fact and, as such, is a subjective judgment. What this means in the quantum case is that to say there is a "click" is one thing: Presumably that is not a subjective judgment if I say it, and Steven says it, and everyone else we talk to says it. However, to say that that means the particular outcome E_b occurred in the POVM $\{E_b\}$ is to lay down a proposition, a subjective judgment.

The reason we could get confused in the classical case, and think that a proposition was more than a subjective judgment is because in the classical case, propositions don't entail probability assignments.

This distinction I'm drawing is not so different than the one Marcus Appleby draws in criticizing the Meyer-Kent-Clifton "nullification of the Kochen-Specker theorem." See his paper, `quant-ph/0109034`.

Schackcosm 41 : *I am afraid that I still don't know precisely where you stand, despite of your effort at explaining.*

Let's start from the classical notion of certainty. Let's consider the case where a physicist is certain that some outcome will not occur. Dutch book consistency implies that the outcome will be in his nullspace. That's the quantum part.

Do you agree?

Yes.

Schackcosm 42 : *1.) Carl is certain that up will occur and Chris is certain that down will occur.*

2.) Carl is certain that there are three chairs in the room, and Chris is certain that there are two chairs in the room.

In both cases their beliefs are contradictory in the same, classical sense.

Do you agree?

No, I do not think the statements have the same meaning. In the first case, in order to find out which of us is "right" and which of us is "wrong" we must elicit the world to produce something that it didn't contain beforehand—namely, the result of the measurement. In the second case, we can go blissfully along thinking that one of us is "right" and one of us is "wrong" simply because the world has something in it that one of the two of our brains is mirroring correctly.

There is a difference. In the classical world, reality is the ultimate arbiter of truth. In the quantum world, where we are fairly convinced that "unperformed measurements have no outcomes," we are actually lucky in a way that there is still an arbiter of agreement—we just can't identify it with a preexisting reality. It seems to me this is a feature of quantum mechanics: The theory can still bring us to agreement even without a preexisting truth value for our propositions. One might have imagined a more malicious world where we would not have even been able to rely upon that.

Schackcosm 43 : *If we make claims, we are COMMITTED to those claims (e.g., via betting behavior implied by the claims). I believe that starting from the notion that different scientists (different agents in the same linguistic community) should not have contradictory beliefs is eminently reasonable. To throw this notion over board, one needs excellent reasons. I enjoy playing with the idea, but I am far from converted.*

This hits upon what I wrote to Carl yesterday. The gulf that separates us seems only to be in where we think this agreement must be applied to get the engine running. I say both of us accepting QUANTUM MECHANICS as a structure is good enough. (I.e., accepting the theory is our common agreement.)

Now to Carl:

Cavesism 49 : *Your point, as we see, is that we can think about life going on after finding a result deemed to be impossible. As you point out, there is a limit (add on the null subspace with epsilon eigenvectors, get a result in the epsilon subspace, update, and then take the epsilon goes to zero limit) in which we can think of updating a state assignment based on outcomes in the null subspace. But I think this misses the point. This isn't updating a prior belief. It's realizing that your prior belief is full of it and abandoning it in favor of life going on, as you put it.*

Moreover, the really nice distinction between classical and quantum Fuchs consistency is lost if we adopt your point of view. If we adopt your way of formulating Fuchs consistency, then it has no content either classically or quantum mechanically.

I'm just repeating myself now, I but I don't see that as contentless at all. In fact, though the effect is the same in both theories, the content is quite different across the two of them. In the classical case, we can always "pick up the pieces" as you say, by realizing that there is something really there and just checking what it is. In the quantum case, we can always bend the world into something we *will* agree upon, even if we violently disagree upon the meaning of some subset of our previous interventions.

Here's the way, Josiah Royce put it in a letter in 1888:

> Thus called upon to explain amid the trade-winds, and under the softly flapping canvas, the mysteries of [quantum mechanics], I put the thing thus: "There was once a countryman," I say, "from Cape Cod, who went to Boston to hear Mark Twain lecture, and to delight his soul with the most mirth-compelling of our humorists. But, as I have heard, when he was in Boston, he was misdirected, so that he heard not Mark Twain, but one of Joseph Cook's Monday Lectures. But he steadfastly believed that he was hearing Mark. So when he went home to Cape Cod, they asked him of Mark Twain's lecture. 'Was it *very* funny?' 'Oh, it was *funny*, yes,—it was *funny*,' replies the countryman cautiously, 'but then, you see, it wasn't so *damned* funny.' Even so, Captain," say I, "I teach at Harvard that the world and the heavens, and the stars are all *real*, but not so *damned* real, you see."

Cavesism 50: *The parties can always come to agreement, no matter what their state assignments, simply by getting amnesia regarding their prior beliefs and then picking up the pieces in the only way they know how.*

But there's really more to the story. They can always come to agreement, indeed—regardless of how disparate their initial opinions—if they are willing to make an essentially infinite expenditure toward laboratory technique. That is to say, the only thing that will give assured agreement in all cases is a set of Kraus operators all of rank-one. Jacobs and I called those infinite strength measurements: the idea being that they are hard to actually do. In more real-world measurements, where the operators are never really rank-one, coming to final agreement will generally require some initial agreement. Whence the point in my August 7 letter to Mermin.

Cavesism 51: *I believe that Fuchs consistency is about coming to agreement in the light of the outcome of an agreed-upon measurement where no party has to abandon his prior beliefs (certainly one has to agree that this is a legitimate case to consider). It could be that all parties are dumbfounded by the result, but let's put that case aside. For all other outcomes, the point is that all the parties be able to come to agreement by updating their prior beliefs. This imposes a strong constraint classically—all parties must have the same support—but appears to be no constraint quantum mechanically. That's an important distinction, it seems to me.*

I do agree that that is a legitimate case to consider. What I am not seeing presently is that its study will shed some foundational light. But I think I'm open-minded on this one: I might be convinced yet; I just don't see it now.

08 October 2001, to Summhammer, "Larger, Smaller"

Thank you for your wonderful, thoughtful, long letter! I have now read it several times, and each time I think I've gotten a little more from it. Thank you also for your concern over my family and associates in light of the September 11 attack: As far as I know, all my friends, and my friends'

friends were left unscathed physically. But it is all a very frightening affair, and it is certainly weighing on everyone's life on this continent and the world.

Concerning the content of your letter, let me especially thank you for the large number of YES's you wrote into the margin of my paper! Let me make a couple of small comments on your one NO.

Summhammerism 3 : *It appears to me that quantum theory is the correct way of reasoning, and classical probability theory is a certain limit of it. But both spring from the SAME way of reasoning. So far, physics has stood in the way of clarifying this. The perennial talk of systems and properties of systems one is forced to carry along when dealing with quantum theory is a real hindrance to clear thought. Remnants of mass points, forces, fields, etc. always sneak in, and with it the need to allude to an objective world out there. As if repeatably detectable structures in the statistics of probabilistic events and their efficient description weren't objective and "out there" enough (to me mathematical truths and the Himalayas are equally "out there"; the former are mastered by acts of mental climbing, the latter by acts of physical climbing, but both require willful action to be conquered. You call it free-standing reality.)*

This issue has now come up in my email a couple of times since the Sweden meeting. Here's the way I put it to Jeff Bub on the last round:

> The main thing is that it sounded like a good opportunity to pound out the similarities and distinctions between our points of view on quantum mechanics without being interrupted every three minutes.
>
> I know I suggested I would write a longer letter soon, but I'm going to wimp out of it again for now. It would concern the main point of distinction I see between us (and also between myself and Pitowsky). Namely, A) that I view a large part of quantum mechanics as merely classical probability theory (which on my view may be an a priori "law of thought") PLUS an extra assumption narrowing down the characteristics of the phenomena to which we happen to be applying it to at the moment, while B) you are more tempted to view quantum mechanics as a *generalization* of classical probability theory (and with it information theory). I know that my view is not fully consistent yet, especially as I have always distrusted mathematical Platonism—which you pointed out to me I am getting oh so close to—but it still feels more right (to me, of course). Ben Schumacher, Ruediger Schack, and I had a long discussion on this (on a long walk) the day after the round table, and I'd like to record that too. Ben took a stance quite similar to yours, and maybe even Ruediger did too (despite his overwhelming Bayesianism). So, I may be the lonely guy out on this. And my view may be subject to change.

To some extent, I can understand both motivations, i.e., to see quantum theory as the larger of the two structures, and alternatively to see it as the smaller of the two. My thoughts are not completely set yet about which direction is the best direction, but let me try to explain a little about what I mean by probability theory possibly being a larger structure than quantum mechanics. (As evidenced in my paper, this is certainly the direction I lean most toward presently.)

Consider some physical system, say my house. And consider some set of questions you might ask about it. For instance, what color is it? (The answers being R, O, Y, G, B, I, V.) Or, what kind of flooring does it have? (The answers being wood, tile, vinyl, or carpet.) On so on: Consider every question you might ask about it. If you were a Bayesian, you would not hesitate taking all the information you know about me and applying it to the construction of a probability function for the outcomes of each such question that could be asked. For instance, if you had gained the

impression in Växjö that I am a sentimentalist, you might place a higher probability on my floors being wooden than otherwise.

However it is also part of the Bayesian creed that there is no such thing as an invalid probability assignment. If there are no logical connections between a set of questions, then there are no constraints on the probability assignments I might give for their potential answers. So, for instance, though you might put a peaked distribution on the answer to the question about my floors, you might put a flat distribution on the colors. And so forth, for every elementary question that might be asked about my house.

However, when we come to quantum mechanics something changes about this. Now, the elementary questions correspond to POVMs. But, using Gleason's theorem, we are no longer free to assign probabilities to their outcomes willy-nilly. All but a very few probabilities assignments are tied together via the existence of a density operator. For instance, viewing quantum probabilities as Bayesian probabilities, one is completely free to assign any probabilities one deems relevant to the outcomes of a σ_x, σ_y, and σ_z measurement. However, once that is done, one is no longer free to specify an arbitrary assignment for spin in the n direction, for any other n.

From this point of view then, quantum mechanics allows only a subset of the vastly larger set of probability assignments one might make to the answers of the physical questions one might ask. And one might think that restriction is accounted for by some physical fact—the yet-to-be-discovered fact that is the essence of quantum mechanics.

Summhammerism 4: *Still, it is in this connection that I wrote a NO into your paper. On p. 28 you say "Probability theory alone is too general of a structure.", and at some other place you say there must be an input from nature. Based on my own games with these questions I doubt this. I think quantum theory is already contained in the basic notions that lead to probability theory. The sum rule of probability is no obstacle, if you ponder what mutually exclusive means from an operational point of view. For this reason I see a valuable contribution in Lucien Hardys attempt of starting from a few axioms, although an axiomatic approach is unsatisfactory as long as the axioms arent simple truths instead of formal assumptions.*

So, indeed, what I said above is what I meant as an input from nature: It is whatever binds us to Gleason's theorem. (Gleason's theorem being the string that ties all the various distributions for a physical system together.)

You'll note actually that Hardy is almost an antithesis to this idea. He starts with structures that are larger than both classical probability theory AND quantum mechanics. By adding an extra postulate he can narrow it down to either one or the other (or any of a number of other structures). What I want is start with classical probability and then narrow it down to quantum mechanics.

It could be the wrong direction, but it is the one that feels predominantly right to me (and the one that seems to me to have the highest probability of leading to some interesting physics). It is a subjective judgment of course, but that's all that each of us has.

By the way, Caves and Schack and I have been thinking about applying for a visit to the Erwin Schroedinger Institute next spring or early next summer. The plan would be to write (the bulk of) an RMP article on all this Bayesian business while we're there. It'd be great to have your ear to test it out on, if we do follow through with the plan.

09 October 2001, to Mermin, "Writing Physics"

I'll place the new supplement to the Samizdat in the next email as plain text. Please let me know whether you're able to TEX it up fine; if that doesn't work out, I can post it on my web page (as a PostScript file) as I did before.

Collecting it up, it's hard to believe I've written this much in the little time since Växjö. I guess it's been an active time for me. I think there's no doubt that I've gone through a phase transition. For all my Bayesian rhetoric in the last few years, I simply had not realized the immense implications of holding fast to the view that "probabilities are subjective degrees of belief." Of course, one way to look at this revelation is that it is a *reductio ad absurdum* for the whole point of view—and that will certainly be the first thing the critics pick up on. But—you wouldn't have guessed less—I'm starting to view it as a godsend. For with this simple train of logic, one can immediately stamp out the potential reality/objectivity of any of a number of terms that might have clouded our vision.

You'll find the most useful stuff in here starting at my last note to you, i.e., page 78 onward. In particular you might enjoy the chart on page 84. It shows, I think, that when this exercise of epistemologizing so much is over and done with, there's still a fair amount left that one might be willing to call concrete reality.

I do hope you get something out of this. Two of your questions were the sources for the vast majority of the pages in the document. If you hadn't pushed me, I may have never seen that so much was waiting in the wings to be made sense of.

16 October 2001, to Waskan, "Craters on the Moon"

In a way you stole me away from the family tonight; the little ride home was full of thoughts about craters on the moon. You said something like, "That there are a thousand craters on the dark side of the moon, is a true statement regardless of whether it's useful or not."

Below is a passage I took from David Darling. It once made a good impression on me (long before my James days), and I can't help but feel that it is relevant to tonight's conversation too. If all the world is but atom and void—or substitute your favorite metaphysic, for that matter—then, it seems to me, there is no strict sense in which there are "craters" on the moon at all. To interpret the coarse information I have about some aspect of my experience as a map of the craters of the moon, seems to me an ultimately subjective judgment—one that I make because it is more or less useful.

Of course, I'm not wedded to these ideas: the game I play is to pick and choose anything from any philosophy that will help me make sense of the physics I'm doing and to promote it to a new level. Sometimes, after a sufficient amount of play, I change my philosophical mind. But I think the observation below is not a completely idiotic one, and it sends me some way toward the pragmatic conception of truth. No proposition I use in my daily life can be strictly true or false in the sense of reflecting the world as it is independently of me. And if not the case in my daily life, then why should it be the case at some more "ultimate" level (i.e., fundamental physics) that—after all is said and done—was intellectually derived from that daily experience in the first place?

> The interface between mathematics and everyday reality appears sharp and immediate at this point: one sheep, one finger, one token; another sheep, another finger, another token, and you can take away tokens or add them, as you can with your fingers. The tokens—the numbers—are just abstracted fingers; the operations for dealing with the tokens are just the abstracted raising or lowering of the fingers. You make a

one-to-one correspondence between the tokens and whatever it is you want to reckon, and then forget about the fingers.

At first, it seems clear from this that mathematics must be somehow already "out there," waiting to be discovered, like the grain of the stone. One sheep add one sheep makes two sheep. Two sheep add two sheep makes four sheep. That is certainly the practical end of the matter as far as the shepherd and the merchant are concerned. But already, even in this most simple mathematical maneuver, something strange has happened. In saying "one sheep add one sheep" we seem to be implying that any two sheep will always be identical. But that is never the case. Physically, the first sheep is never exactly equal to the second: it may be a different size, have different markings. It takes only one molecule to be out of place between the two, and they are not identical. Indeed, because they are in different places they are inevitably not the same on that basis alone. We have extracted a perceived quality to do with the sheep—namely, their "oneness," their apartness—and then merged this quality by means of another abstraction—the process of addition. What does it mean, physically, to "add" things? To put them together? But then what is "putting together" two sheep? Placing side by side, in the same field—what?

All this may seem like nit-picking. But on the contrary, it brings us back to the central mystery—the relationship between the inner and the outer, the world of the rational mind and the world "out there." In the physical world, no two sheep are alike. But, more fundamentally, *there are no "sheep."* There are only some signals reaching the senses, which the left brain combines and then projects as the illusion of a solid, relatively permanent thing we call a sheep.

Like all objects, sheep are fictions: chimeras of the mind. It is our left hemispheres, having through natural selection evolved this skill for extracting survival-related pieces of the pattern, that trick us into seeing sheep, trees, human beings, and all the rest of our neatly compartmentalized world. We seek out stability with our reasoning consciousness, and ignore flux. We shut our eyes to the continuous succession of events if those events seem not to substantially affect the integrity of what we see. So, through this classifying and simplifying approach we make sections through the stream of change, and we call these sections "things." And yet a sheep is not a sheep. It is a temporary aggregation of subatomic particles in constant motion—particles which were once scattered across an interstellar cloud, and each of which remains within the process that is the sheep for only a brief period of time. That is the actual, irrefutable case.

17 October 2001, to Waskan, "Quick Single Point"

Thanks for the note: I'm still digesting it. But let me quickly reply to the one thing that I can reply to.

Waskanism 1: *Also, it seems strange for the fellow you quote to concede that there is light, eyes, the left half of one's brain, a process known as natural selection, and at the same time to deny the existence of sheep.*

Indeed, it seems strange to me too: He's pretty clearly not being consistent. But the role of the quote for me was as a motivating piece. (I read Darling a few years ago, and the main reason I used the quote last night was because it was already in my computer and I could, thus, send off

a quick note to you.) My sentences above the quote were meant to show that I toy with the idea of going a further (more consistent) step:

> And if not the case in my daily life, then why should it be the case at some more "ultimate" level (i.e., fundamental physics) that—after all is said and done—was intellectually derived from that daily experience in the first place?

It is going that extra step that seems to me to be heading down the track to Jamesianism. Let me read your note again ...

17 October 2001, to Waskan, "Quick Second Point"

Sorry, I can't wait for the beer.

Waskanism 2: *Here's a daily affirmation for you. Look in the mirror (not Rorty's mirror of nature, the one in your house), and say the following:*

I am pretty darned sure that I exist. I think there is other stuff in the universe too. I'm pretty darned sure that I have beliefs about (i.e., representations of) the other stuff that might be out there. My beliefs are true insofar as the world is how I represent it to be. If the world outside of my mind is in no way how I represent it as being (e.g., IF THERE ARE NO SHEEP, etc.), then all of my beliefs are false. Even if all of my beliefs about the things outside of my mind are false, this should have no bearing on the nature of truth itself (i.e., correspondence).

If you want to use true and false that way, then—in my present state of mind—yes, I would say that, strictly speaking, most all my beliefs are false. Beliefs play the role of coordinating our actions, and, in that way, can be more or less useful. But (in the too small of thought I've given this) I can't find a role for the concept of belief outside its use.

Like you, I am pretty darned sure of the existence of a world outside myself. But I would say that that surety comes NOT from some (transcendental?) knowledge that my beliefs mirror that world as it is. It is just the opposite. I believe in a real world outside myself because, throughout my life, things continue to take me by surprise. Significant numbers of my beliefs are systematically INvalidated with each new day. There's my evidence of the real world. Below is the way I put this point in a recent paper.

The point of view is not completely worked out yet—and it may never be—but my experience in quantum mechanics makes it feel more right than the other options I've seen so far.

Knowing me, I'll probably give your note another read, and be back again tomorrow. I hope you don't lose patience with me.

17 October 2001, to Waskan, "Quick Third Point"

OK, one more for the day.

Waskanism 3: *Also, sure no two sheep are exactly alike. Neither are any two bachelors. That doesn't entail that there are no bachelors. X is a bachelor if X is unmarried, and X is male, and X is eligible (e.g., not a priest).*

Granted. But what I thought was at issue is whether there are "bachelors" without the agents who make up (and use) all these judgmental categories. If all the world is BUT atom and void—to

use an allegory I like but which should not be taken literally as my view—then where do all these extra distinctions come from if not the judgmental agent?

I say I thought this was the issue because my reading of the pragmatic conception of truth is more to the following point: Without agents, there are no "propositions." Therefore "propositions" cannot be true or false in any absolute sense. Without the agent, there is the world, and it is just whatever it is. A proposition adds something to the world that it itself did not possess before the agent's attention was drawn to it (via the act of dreaming it up, writing it down, acting upon it, etc.).

There's probably nothing worse than to have an armchair philosopher in your presence ...

18 October 2001, to Waskan, "No Doobies Here?"

Waskanism 4: *Lots of folks want to say that there are joints in the natural order (though, admittedly, bachelorhood probably isn't one of those joints).*

I'm not sure what you mean by the term "joints." Can you define it precisely? (Not knowing what you mean, it leaves me unable to reply to most of your message.)

Waskanism 5: *If, however, you take a step back and look at the big picture, ... I suspect that the same can be said for the relationship between atoms and tables.*

You know what my worry is (fueled by the 75 year debate on the interpretation of quantum mechanics). It's that we just can't step far enough back. We are immersed in this thing called existence, and there's just no way to get a view from outside it. We do the best we can from the inside, and that's called science. (For me, the phrase "best we can" means to eliminate impredictability—i.e., (only half jokingly) to delete reality as much as possible. Cf. yesterday's note about Emma.)

Tell me what a joint is, so I can think a little more about what you said.

22 October 2001, to Schack-Caves, "The Dilemmas of Subjectivism"

I apologize for holding off so long in a reply to Ruediger's letter concerning the RMP article. The difficulty has come in that I didn't know how to reply. (I guess I still don't.)

The point of some potential consternation is this:

Schackcosm 44: *What I think we should be doing is a paper on "Interpretations of probability in quantum mechanics (with special emphasis on the Bayesian viewpoint)". The paper would NOT be on the interpretation of quantum mechanics.*

The problem is, I don't see how to separate the two issues. Where does the interpretation of probability fall off and the interpretation of quantum mechanics kick in? How can one have an interpretation of quantum mechanics wherein the wavefunction is objective, but still think of probabilities as being strictly subjective? Similarly, vice versa?

Let me do this: Let me ship to you both the mini-samizdat of my thoughts that came out of my post-Vaxjo broodings. You tell me which *sentiments* will be banned and which won't if we end up skinny-dipping together. I can foresee some being excluded—like the stuff in my letter to Wootters—and I can accept that; but for the greater majority of the writings, I can't see myself drawing a line, and I'm wondering where you will draw it.

Looking back on the BFM debate, I think the most important thing to come out of it for us three in particular is that it makes it absolutely clear that we need to get our thoughts straight on the "principal principle" before we can embark on a consistent statement of our position. For I see no way to erase the dilemma: Either we accept that the ascription of one Kraus rule over another in a measurement intervention is a subjective judgment, or we accept a quantumatized version of Lewis's principal principle. Why did we all reject the principle before if we find ourselves accepting it now? This is something we ought to reflect upon more deeply.

23 October 2001, to Schack, "United We Stand Airlines"

Schackcosm 45 : *Do you think you could give me a lift at such an early time?*

I do whatever it takes for the greater good of quantum mechanics. Of course I can give you a lift! (If you can stand to listen to me babble that early in the morning.)

26 October 2001, to Butterfield, "The Feynman Cult"

I don't remember railing against the Feynman cult in your presence, but your letter gives some evidence that I must have. (Doing such things was a common pastime for me at Caltech, but I've mellowed a little in my efforts since leaving, i.e., since the cult hasn't been in my face on a daily basis.)

But, anyway, yes I told Brandt that I would come. (I hope you will come too.) I've even thought about how I will open my talk: with the Feynman quote below. Lord knows I'm no materialist, so you can rest assured that I'll do my best to "zing" it up afterward. (See my Samizdat, page 237.)

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generation of creatures, what statement would contain the most information in the fewest words? I believe it is the atomic hypothesis (or the atomic fact) that all things are made of atoms—little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into on another.
>
> Everything is made of atoms. That is the key hypothesis.

27 October 2001, to Caves-Schack, "Re: Coming to Agreement"

Schackcosm 46 : *In any case, do you know a reference for strong Dutch-book consistency?*

Nothing comes to mind: I think I first learned of it from you (was it in Cambridge?). Might it have been in either of those papers we looked up in a trip to the university library one day? This was one of them (though I don't have it anymore):

> J. G. Kemeny, "Fair Bets and Inductive Probabilities," J. Symb. Logic **20**, 263–273 (1955).

Ahh, here was the other one:

> R. S. Lehman, "On Confirmation and Rational Betting," J. Symb. Logic **20**, 251–262 (1955).

It might be a good idea to dig those up again.

I'll try to answer your other questions next week, after I get a chance to get a better grasp on the paper.

27 October 2001, to Schack, "Literature"

Schackcosm 47: *A question to you: How much literature is there on interpretations of probability in q.m.? Would it be feasible to review it all?*

Attached is everything I had collected up previous to the Cerro Grande fire. A skim of that document might answer your question in the most direct way.

By the way, I've asked Maria Carla Galavotti for her criticisms (and any others that she knows) of the principal principle. But I haven't gotten a reply from her yet.

Thinking back on it, I can't remember if we ever had any NONquantum reasons that were substantially different than J. S. Mill's argument against a "substance" underlying the phenomenal world:

> If there be such a *substratum*, suppose it at this instant miraculously annihilated, and let the sensations continue to occur in the same order, and how would the *substratum* be missed? By what signs should we be able to discover that its existence had terminated? Should we not have as much reason to believe that it still existed as we now have? And if we should not then be warranted in believing it, how can we be so now?

Similarly one could say of objective probability: Bayesian coming-to-agreement would work precisely the same whether the objective probability is there or not. But did we have any other arguments than that (that did not depend on upon quantum mechanics, for instance in the nonuniqueness of the density operator decomposition)? Did we have any examples where believing in objective probability in Lewis's sense would be downright misleading in how it might suggest tackling a practical problem?

I so wish I had my old file cabinet back again; writing you this morning has made that feeling more acute.

04 November 2001, to Peres, "Dreams of an Ever-Evolving Theory"

Asherism 1: *I am reading a wonderful book "Dreams of a Final Theory" by Steve Weinberg. Chapter 7 is "Against Philosophy" and I highly recommend it. I got that book for $7.00+tax in a used books shop in Santa Barbara, where Aviva was looking for something else.*

Thank you for the tip; I did read Weinberg's chapter 7. He writes in a very crisp and no nonsense way, and I like that. (You know I constantly fight tendencies in the opposite direction in my own writing.) His points are well taken, especially the ones about how a preset philosophy can create immense blinders for the scientist:

> "... in rejecting it the [PHILOSOPHY-X]ists were making the worst sort of mistake a scientist can make: not recognizing success when it happens." — page 177, paperback version.

But, I think a deeper point is the one he makes near the beginning of the essay:

> "I do not want to draw the lesson here that physics is best done without preconceptions. At any one moment there are so many things that might be done, so many accepted principles that might be challenged, that without some guidance from our preconceptions one could do nothing at all." — page 167

And I agree with the numerical tally of the next sentence:

> "It is just that philosophical principles have not generally provided us with the right preconceptions." — page 167

However, I part company with him in thinking that that is a strong argument against pursuing philosophy as a sideline to science. That is, I don't know what the preconceptions can be if they're not philosophies.

When it comes to philosophies and, not unrelated to that, scientific research directions, I tend to take a lot of stock in a Darwinian kind of conception. That is, we each should do precisely what we feel compelled to do; we each should research precisely what we feel compelled to research. There's probably nothing we can do about it anyway. Indeed ninety-nine percent of the time we will be on the wrong track: The world supplies a selection pressure for our thoughts, just as it does for the lifespan of the drosophila. And just as it is not possible for the drosophila to change its genetic makeup before it meets its demise, I think the only thing we can do as scientists is cultivate to the best of our ability the philosophical preconceptions that led us down our own paths. Ninety-nine percent of us will be forgotten from the history books, but the ones of us that remain will do so because the world is such that it is less likely that we should fall.

For myself, I have DISCOVERED that I have chosen a direction of thought that is very closely aligned with the philosophical movement of pragmatism from the early part of the 20th century—a movement the details of which have been nearly forgotten in modern times. Interestingly, the thing that set me on to this realization was Martin Gardner's essay "Why I Am Not a Pragmatist" in his book *The Whys of a Philosophical Scrivener*. (You probably remember Gardner from his column in *Scientific American*.) This happened about three months ago. I really recommend you read the article if you get a chance. Maybe your library has a copy of the book. I think in reading it, you will discover that the essay might just as well have been titled, "Why I Am Not in Agreement with Fuchs and Peres's Physics Today Article." For, with each reason Gardner used to explain why he was *not* a pragmatist, I found myself thinking of quantum mechanics and saying to myself, "ahh, I guess that means I *am* a pragmatist." Really, the analogy is *that* close even though the article has nothing to do with quantum mechanics per se.

The issue is no less than whether "unperformed measurements have no outcomes." The pragmatists, for various reasons, thought it was *safer* to assume that they didn't. The movement then spent the greater part of its time developing the (liberating) consequences of this supposition. Gardner (and Bertrand Russell and G. E. Moore and gazillions of others) thought "how silly" and "how contrived" when it is so much easier to use standard realist language to describe the outcomes of experiments—to assume the outcomes are there before one has a look. But you and I know better, of course. And, I think it is quite useful to know that there was a set of people carrying through the detailed consequences of this line of thought for their broader worldview long before you and I were on the scene. The way I view it, these old thoughts can be a resource to our explorations of quantum mechanics just as much as any other: However, their use is in setting the directions for potentially fruitful lines of thought ... but that should always be the use of any philosophy for any scientist.

Below I'll attach a letter I wrote Bill Wootters a while ago on a similar subject. What I write just after the quote of William James better explains why I chose the title that I did for this note.

19 November 2001, to Caves-Schack, "A Lot of the Same"

Schackcosm 48 : *Before you have me burned alive, please tell me why I am wrong!*

Come on, you know it's the Thanksgiving season. I would never burn you, only roast you. (Though my brother-in-law once fried his turkey in hot oil.)

Sorry for the hiatus, but I just got inundated with email last week, and I didn't have the proper mentality for replying to any of it. So I shut down for a while. Now I'm stuck with trying to clean out an even bigger pile of old mail. But let me compliment you by letting you know that I'm tackling your letter first! (It's the only interesting one in the lot.)

Schackcosm 49 : *I started writing up a summary of our discussions, and hit upon a difficulty when I tried to formulate exchangeability for models. Here is the problem.*

In the traditional formulation of exchangeability, we say that we have N identical systems (same Hilbert space). At this stage, it is thinkable to assign a different state to each system. We then make the judgement that the state of the N systems is exchangeable.

For models, we say that we have N identical apparatuses. Alternatively, we say that we use the same apparatus to perform N measurements, let's say on different, independently prepared systems. At this stage, I can't think of a good reason to even consider assigning different models to the apparatuses. It seems to me that one is forced to say that each apparatus performs the same operation, so should be described by the same model. That leads immediately to a heresy: there exists a true model. If we don't know it, we assign probabilities to models. What we wrote down on your whiteboard is consistent with this viewpoint. The difficulty we encountered formulating exchangeability could mean that it is an unnatural concept in this case. Writing down a mixture of N-fold products of models is completely natural however.

Let me try to allay your fears. I think the issues here are almost precisely the SAME as they are in our old de Finetti considerations. To say it in a way that maybe Carl would endorse, "It's really all about learning." Or in a way that I'm more tempted to these days, it's all about demonstrating a willingness to update one's beliefs—one's commitments, one's pragmatic strategies for action, one's betting behavior—in the light of factual data.

Let me start with an example that's essentially already well-worn for us by now. Suppose we have a rather complicated quantum measurement device whose manufacturer purports it to be the best σ_z measurement device ever built. Furthermore, suppose we have a fresh supply of 10^8 calcium atoms, all meticulously prepared to have spin-up in the x-direction. What do we expect to happen if we individually dump all the atoms into the measuring device? We expect about 50% of them to get registered as spin down and about 50% of them to get registered as spin-up. But what happens if one after another, all the registrations are of the spin-up variety? Well, that outcome sequence would be no less likely or no more likely than any other outcome sequence if we walked into the laboratory with such a radically adamant prior belief. In a real-life situation, however, we would be shocked; we would update our beliefs accordingly—for we would have allowed for the possibility of "learning."

But in this situation, notice that there are at least two extreme cases to which we could attach the possibility of learning. The learning could be about the device or it could be about the preparations. Who's to say that the learning is about something more objective in the one case than

the other? Prosaically, it takes both ingredients (the preparations and the device) to certify the device, and you can't get away from that.

Let me try to tighten this up by sketching how we ought to start thinking about a de Finetti theorem for unknown quantum models. I run a measurement device on N independently and identically prepared quantum systems. Suppose I am absolutely confident of these preparations—i.e., with respect to them, there is nothing left to learn in the technical sense of i.i.d. statistics for any repeated and KNOWN measurement. Then, what can it mean—from a Bayesian point of view—that the measurement device works according to an unknown model? It means that after all the outcomes are gathered, there's still something left to be learned from the posteriori quantum state for the systems that were measured.

That is, more simply, the best judgment we can make about the systems that passed through the measurement device is that they are exchangeable CONDITIONED on the registered measurement outcomes. For instance, suppose the device spits out an index i at each round. The issue is, what pragmatic meaning should we give to each such i? Quantum mechanically, the predictive meaning of an index i is specified by the Kraus operator A_i we associate with that outcome. (It's retrodictive meaning is given by the positive part of A_i—the POVM.)

If we think we don't exactly know what the device is doing to each individual system, then we shouldn't yet dare to make an association $i \to A_i$. (To make an extreme point of it, for all we know, the device might be entangling all our test systems.) We should just rest confident that no matter what order we send the systems through the device, we will end up with the same subjective beliefs in the end. Thus, if we gather up all the systems for which an outcome i occurred (as opposed to some other outcome j), then the subjective density operator we assign should be exchangeable. Using the standard quantum de Finetti theorem, we then get that that density operator must be of de Finetti form. Writing each of the final density operators as a linear map acting on the initial density operator, we (should) get something like the desired theorem for unknown quantum models. If we believe that we can learn something about the model, then the probability distribution that appears in the de Finetti form is restricted to being something other than a *delta*-function.

In summary, our belief that the best we can say of the outputs is that they ought to be exchangeable (conditioned on the factual outcomes), leads directly to a notion of mixture of models—i.e., that the output density operator is controlled either by a Kraus operator A_i, or a Kraus operator B_i, etc., etc., about which we capture our ignorance through some subjective probability distribution.

Now, just as the regular de Finetti theorem cannot put an end to the principal principle, we cannot use this (proposed) theorem to put a stake through the heart of the true believer of objective quantum models. That is, David Lewis might say of the regular de Finetti theorem, "That is a very nice theorem, but it doesn't change the fact that there really is always a 'man in the box.' His name is God." And so he would probably also say of our quantum models (if he knew quantum mechanics). Instead, all the (proposed) theorem can do is show that it is *possible* to get by without a man in the box. We don't need him; all we need is something like the judgment of exchangeability for the outputs (conditioned on the outcomes) along with i.i.d. on the inputs.

So that's the sketch. Now, how to dot the i's and cross the t's? I can foresee at least one difficulty that I'm not clear-headed about right now. That is, using the description above, for each index i, we will generate a probability distribution over models. But by what regularity condition can we assure that $p_i(\mathcal{A}) = p_j(\mathcal{A})$? What I mean by this notation is that $\mathcal{A}$ stands for the model in total (i.e., all the Kraus operators in it) and $p_i(\mathcal{A})$ stands for the probability distribution in the de Finetti theorem derived for each index i. It is probably so simple as this: If we were to imagine doing tomography on the posterior states for each index, then the states derived from that should always average up to a valid density operator. But I'm not exactly sure how to put that idea into action.

Oh, and here's another intriguing point that ought to be explored. Suppose we focus our attention on a given exchangeable density operator. There are many ways that operator could arise, but suppose that it came about as the posterior state arising from many identical measurements (in the sense above). The question I have in mind is how much freedom do we have for trading off between an unknown preparation and an unknown model for getting to the final state? Can one always find a fixed initial preparation and a mixture of models that will give rise to the final exchangeable state? Can one always find a mixed model and a mixture of initial preparations? Probably yes and yes, but I'm not completely sure.

I'm so glad to hear that you may have reversed your opinion on William James (at least a little). By the way, I hope you notice how these de Finetti considerations are drawing out a lot of the considerations I was trying to express to you in Samizdat-II and our subsequent discussions. For it helps draw the distinction between the amorphous index i that arises in a quantum measurement and the meaning A_i that I ascribe to that event. The symbol A_i plays the role of a proposition that I write about i: It carries the information about how I will react after having seen it, how I will place my bets.

I'm willing to believe this whole debate about "truth" might be a red herring—i.e., that we might easily be able to get away with never uttering the word. But I think the realization of the last paragraph had a primitive expression in James's worries about "truth" nevertheless, and to that extent maybe he and his movement of pragmatism are worth contemplating (though of course not subscribed to in toto).

19 November 2001, to Schack, "Re: William James"

Schackcosm 50 : *It's a curious mix, though. Sometimes he seems to be very close to Kantian ideas on truth, then he seems to subscribe to a naive correspondence theory of truth, at least for simple facts such as "this detector has clicked."*

Yes, I'm more attuned to that now, and I'm trying to get it straight through extended readings (such as A. J. Ayer's book *The Origins of Pragmatism* that I picked up in New York City with you).

Schackcosm 51 : *I am also a little disturbed by his praise for Ostwald.*

I don't remember his praise for Ostwald. Who was Ostwald? And what does James praise him about?

Schackcosm 52 : *I am half way through writing a much improved draft of our paper.*

Sounds great.

20 November 2001, to Caves-Schack, "A Bathtub Moment"

I'm sure I've already told you both the story of the time I happened to end up at a British pub with Caroline Thompson, the famous Bell-inequality conspiracy theorist, but let me repeat it for the purpose of having it in this box. Somebody at the table was speaking of the great importance of intuition, of being able to "see beauty in a theory." I, with my usual example, piped up that I thought that was hogwash: I always thought Mary Ann was the prettiest girl on Gilligan's Island; my best friend thought it was Ginger. Anyway, I followed that comment with, "I never use any

intuition in my calculations; I don't even know what intuition can be in that context." Caroline Thompson harrumphed, "Well we could see that from your talk!"

OK, so it won't be an intuition, but here's a hunch that hit me while I was taking a shower this morning. It's connected to the long note on de Finetti I sent you yesterday.

Among the thirty other reasons I have been thinking that trace-preserving quantum operations (and now measurement models explicitly) are subjective entities is because one can make a one-to-one correspondence between them and the density operators on a larger Hilbert space. That is, they have the same structure as the states of belief that we've already toyed with.[1] This has suggested to me that there ought to be a Gleason-like theorem for quantum operations (which I pursued a little bit but never could quite make things connect). But now this idea is rearing its head again in the context of a de Finetti theorem for models.

Yesterday, I blithely said something to the following effect. To get at a concept of an unknown model, what you do is 1) to 4) :

1. Start with many copies of a quantum system, for which you believe 1) that they are exchangeable and 2) that there is nothing left to learn. The standard quantum de Finetti theorem then gives us that the density operator we ascribe to the collective system will be a tensor product of identical quantum states.

2. Now drop each of those systems into a measurement device and note the outcomes i. Separate the post-measurement systems into bins according to those outcomes.

3. Finally suppose we believe that the quantum state we ascribe to each bin ought to be an exchangeable state for which we *can* learn something. The standard de Finetti theorem gives that this state must be of de Finetti form (with a nontrivial support).

4. The hunch was that the conjunction of 1), 2), and 3)—or them along with some minor additional regularity condition—would specify the content of the phrase "an *unknown* quantum measurement model." The unknown model is simply given by making explicit the form of the linear maps connecting 1) to 3) for all possible inputs into A.

Notice that nowhere in there did I say anything about these maps being completely positive. I just chose the word "linear" for some reason. But surely the assumption of complete positivity must come into this too. So now my question: UNDER THE ASSUMPTION that a measurement model is a state of belief, can one adequately explain the notion of an unknown measurement model by de Finetti techniques WITHOUT the technical assumption that these linear maps are completely positive? Or, instead, is complete positivity absolutely crucial to the program?

It strikes me that Carl's superoperator calculus has got to be the way to go for exploring this question.

20 November 2001, to Schack, "James's Loose Lips"

Schackcosm 53 : *In lecture 2, James says "I found a few years ago that Ostwald, the illustrious Leipzig chemist, had been making perfectly distinct use of the principle of pragmatism [. . .]"*

In my own words, Ostwald rejected as meaningless any statement that did not have observable consequences. A very pragmatic attitude indeed. The trouble was that Boltzmann's ideas about

[1] By the way, Carl, as I recall, was never happy with thinking this point had any significance for our program. I think he saw it as little more than a coincidence.

atoms fell into this category, at least that was the public opinion, led by Ostwald, at the time. No cash value in the "atom hypothesis". I imagine somebody like Ignacio (I pick him only because of his obvious scepticism at my talk in Benasque) to ask: where is the cash value in the Bayesian approach to quantum mechanics?

I agree with you now that that is a troubling praise coming from James. I noticed similarly somewhere else in the book that he also classified Mach in the ranks of the pragmatists.

But there is a grave distinction between positivism (Mach, Ostwald, etc.) and pragmatism (James, Dewey, etc.) as I see it. The positivists eschewed all metaphysical assumptions—thus the egg on their faces for not coming up to speed on the statistical mechanics an atomic hypothesis can afford. The pragmatists, on the other hand, are willing to glorify any metaphysics with a cash value. This relates to the passage by James on Mill that I read you while you were visiting. With metaphysics, the cash value is not in its explanation of any previously discovered facts, but in the concrete actions its BELIEF will give rise to in the agent believing it. Thus James's argument, for instance, for everyone's right to believe in a God, even if that God will never have the opportunity of being confirmed or falsified in an objective fashion. A God's validity in an agent is in his cash value for the agent's ethics, morals, and mode of action for his daily life.

To put this in concrete terms for Carl versus me: I would say that the ontological hypothesis I'm shooting for will show some cash value in the amount of interesting physics it leads to, to the opening up of new quantum computing and quantum control and quantum cryptography methods. And I think (or, more accurately, BELIEVE) it will help us make the leap to the next stage of physics. Whereas I would say—but it's just a gut feeling—that the ontology Carl has been shooting for (i.e., the Hamiltonian) has no such cash value. Only the money flow in the banks will ultimately tell. (And unfortunately, that can only be done with hindsight.)

James—I think in his essay "The Sentiment of Rationality"—has a beautifully worded passage on these considerations that I'll try to get scanned in tonight and sent to you. But I think James himself is either not consistent in his writings, or he's pretty sloppy in reading the other writers he wants to praise. (What I know of him now, I think it's probably mostly the latter.)

Sorry for writing all this. I got carried away. I hope to get your new draft printed out today, and studied partially tonight.

By the way, I think Ignacio is a good benchmark with all this. If we can't find a way to impact him in five or seven years, say, then maybe indeed all this is for naught.

20 November 2001, to Renes, "One Horse's Mouth"

I finally get to the last note I owe you.

Renesism 1: *you make this point in the nato paper explicitly (that the unitary taking the quantum bayes rule post measurement state to the orthodox post measurement state is a "mental readjustment" and does depend on the input state). there's a lot here, though, especially since carl initially balked pretty hard at the idea since "we're physicists so i don't know what he's (you) talking about."*

I'm intrigued by your phrase "there's a lot here, though." I don't quite understand what you're trying to get at—that you agree that it's a difficulty, or that it's a good thing? Or that if Carl balked, that might be an unintentional mark in its favor? Can you explain a little better?

Renesism 2: *some days i don't feel like a physicist (happily coinciding much of the time with the days i don't want to) so i'm not initially troubled by this. however, as i said, there's a lot of "stones*

unturned" here. are you saying that there is no physical picture of what's going on, there must be some subjective element "uncaptureable" by a physical picture? this seems to fit with your rep as being an "extreme subjectivist" but i'd rather hear it from the horse's mouth.

I'm not sure exactly how I should reply to this. You probably know my thoughts at this point better than I know them myself. Maybe I should say it like this: My pet idea at the moment is that there was a world here before humankind ever appeared on the scene; there'll be a world here after we disappear. But I would say the world is still under construction; there is no sense and no ultimate level at which it is already complete. To that extent, I believe our beliefs, our passions, our actions, our inventions, and our dreams modify the world and form part of its construction in a nonnegligible way.

And I think our greatest hint of that comes from quantum mechanics. I would say that what we're learning in a precise way from it is that there is something about the stuff of the world that makes it uncaptureable with a purely physical picture. We find that we cannot even draw a picture of the world without including our beliefs and belief changes as a crucial background in the sketch. (How could we if the world's not completed yet?)

Does that make me an extreme subjectivist? I don't know. Whatever it is though that I should be called, I think this willingness to accept a substantial part of quantum mechanics as simply "law of thought" will keep me from going down a misguided path. I.e., the path of trying to ascribe all the easiest terms in the theory a kind of physical reality independent of our presence as active agents.

One horse's mouth.

21 November 2001, to Schack, "Pragmatism versus Positivism"

Both quotes are taken from William James's essay "The Sentiment of Rationality."

Quote I:

> The necessity of faith as an ingredient in our mental attitude is strongly insisted on by the scientific philosophers of the present day; but by a singularly arbitrary caprice they say that it is only legitimate when used in the interests of one particular proposition—the proposition, namely, that the course of nature is uniform. That nature will follow tomorrow the same laws that she follows today is, they all admit, a truth which no man can *know*; but in the interests of cognition as well as of action we must postulate or assume it. As Helmholtz says: "*Hier gilt nur der eine Rat: vertraue und handle!*" And Professor Bain urges: "Our only error is in proposing to give any reason or justification of the postulate, or to treat it as otherwise than begged at the very outset."
>
> With regard to all other possible truths, however, a number of our most influential contemporaries think that an attitude of faith is not only illogical but shameful. Faith in a religious dogma for which there is no outward proof, but which we are tempted to postulate for our emotional interests, just as we postulate the uniformity of nature for our intellectual interests, is branded by Professor Huxley as "the lowest depth of immorality." Citations of this kind from leaders of the modern *Aufklärung* might be multiplied almost indefinitely. Take Professor Cliffords article on the "Ethics of Belief." He calls it "guilt" and "sin" to believe even the truth without "scientific evidence." But what is the use of being a genius, unless *with the same scientific evidence* as other men,

one can reach more truth than they? Why does Clifford fearlessly proclaim his belief in the conscious-automaton theory, although the "proofs" before him are the same which make Mr. Lewes reject it? Why does he believe in primordial units of "mind-stuff" on evidence which would seem quite worthless to Professor Bain? Simply because, like every human being of the slightest mental originality, he is peculiarly sensitive to evidence that bears in some one direction. It is utterly hopeless to try to exorcise such sensitiveness by calling it the disturbing subjective factor, and branding it as the root of all evil. "Subjective" be it called! and "disturbing" to those whom it foils! But if it helps those who, as Cicero says, "*vim naturae magis sentiunt*," it is good and not evil. Pretend what we may, the whole man within us is at work when we form our philosophical opinions. Intellect, will, taste, and passion co-operate just as they do in practical affairs; and lucky it is if the passion be not something as petty as a love of personal conquest over the philosopher across the way. The absurd abstraction of an intellect verbally formulating all its evidence and carefully estimating the probability thereof by a vulgar fraction by the size of whose denominator and numerator alone it is swayed, is ideally as inept as it is actually impossible. It is almost incredible that men who are themselves working philosophers should pretend that any philosophy can be, or ever has been, constructed without the help of personal preference, belief, or divination. How have they succeeded in so stultifying their sense for the living facts of human nature as not to perceive that every philosopher, or man of science either, whose initiative counts for anything in the evolution of thought, has taken his stand on a sort of dumb conviction that the truth must lie in one direction rather than another, and a sort of preliminary assurance that his notion can be made to work; and has borne his best fruit in trying to make it work? These mental instincts in different men are the spontaneous variations upon which the intellectual struggle for existence is based. The fittest conceptions survive, and with them the names of their champions shining to all futurity.

The coil is about us, struggle as we may. The only escape from faith is mental nullity. What we enjoy most in a Huxley or a Clifford is not the professor with his learning, but the human personality ready to go in for what it feels to be right, in spite of all appearances. The concrete man has but one interest—to be right. That for him is the art of all arts, and all means are fair which help him to it. Naked he is flung into the world, and between him and nature there are no rules of civilized warfare. The rules of the scientific game, burdens of proof, presumptions, *experimenta crucis*, complete inductions, and the like, are only binding on those who enter that game. As a matter of fact we all more or less do enter it, because it helps us to our end. But if the means presume to frustrate the end and call us cheats for being right in advance of their slow aid, by guesswork or by hook or crook, what shall we say of them? Were all of Clifford's works, except the *Ethics of Belief*, forgotten, he might well figure in future treatises on psychology in place of the somewhat threadbare instance of the miser who has been led by the association of ideas to prefer his gold to all the goods he might buy therewith.

In short, if I am born with such a superior general reaction to evidence that I can guess right and act accordingly, and gain all that comes of right action, while my less gifted neighbor (paralyzed by his scruples and waiting for more evidence which he dares not anticipate, much as he longs to) still stands shivering on the brink, by what law shall I be forbidden to reap the advantages of my superior native sensitiveness? Of course I yield to my belief in such a case as this or distrust it, alike at my peril, just as I do in any of the great practical decisions of life. If my inborn faculties are good,

I am a prophet; if poor, I am a failure: nature spews me out of her mouth, and there is an end to me. In the total game of life we stake our persons all the while; and if in its theoretic part our persons will help us to a conclusion, surely we should also stake them here, however inarticulate they may be.

Quote II:

Now, I wish to show what to my knowledge has never been clearly pointed out, that belief (as measured by action) not only does and must continually outstrip scientific evidence, but that there is a certain class of truths of whose reality belief is a factor as well as a confessor; and that as regards this class of truths faith is not only licit and pertinent, but essential and indispensable. The truths cannot become true till our faith has made them so.

Suppose, for example, that I am climbing in the Alps, and have had the ill-luck to work myself into a position from which the only escape is by a terrible leap. Being without similar experience, I have no evidence of my ability to perform it successfully; but hope and confidence in myself make me sure I shall not miss my aim, and nerve my feet to execute what without those subjective emotions would perhaps have been impossible. But suppose that, on the contrary, the emotions of fear and mistrust preponderate; or suppose that, having just read the *Ethics of Belief*, I feel it would be sinful to act upon an assumption unverified by previous experience—why, then I shall hesitate so long that at last, exhausted and trembling, and launching myself in a moment of despair, I miss my foothold and roll into the abyss. In this case (and it is one of an immense class) the part of wisdom clearly is to believe what one desires; for the belief is one of the indispensable preliminary conditions of the realization of its object. *There are then cases where faith creates its own verification.* Believe, and you shall be right, for you shall save yourself; doubt, and you shall again be right, for you shall perish. The only difference is that to believe is greatly to your advantage.

The future movements of the stars or the facts of past history are determined now once for all, whether I like them or not. They are given irrespective of my wishes, and in all that concerns truths like these subjective preference should have no part; it can only obscure the judgment. But in every fact into which there enters an element of personal contribution on my part, as soon as this personal contribution demands a certain degree of subjective energy which, in its turn, calls for a certain amount of faith in the result—so that, after all, the future fact is conditioned by my present faith in it—how trebly asinine would it be for me to deny myself the use of the subjective method, the method of belief based on desire!

In every proposition whose bearing is universal (and such are all the propositions of philosophy), the acts of the subject and their consequences throughout eternity should be included in the formula. If M represent the entire world minus the reaction of the thinker upon it, and if $M + x$ represent the absolutely total matter of philosophic propositions (x standing for the thinker's reaction and its results)—what would be a universal truth if the term x were of one complexion, might become egregious error if x altered its character. Let it not be said that x is too infinitesimal a component to change the character of the immense whole in which it lies imbedded. Everything depends on the point of view of the philosophic proposition in question. If we have to define the universe from the point of view of sensibility, the critical material for our judgment lies

in the animal kingdom, insignificant as that is, quantitatively considered. The moral definition of the world may depend on phenomena more restricted still in range. In short, many a long phrase may have its sense reversed by the addition of three letters, *n-o-t*; many a monstrous mass have its unstable equilibrium discharged one way or the other by a feather weight that falls.

26 November 2001, to Caves-Schack, "PRA Proofs"

I'm in the office again finally, and I've read over the PRA proofs. Of course, as always, I can't see that any of their changes were for the betterment of the paper ... but this time at least, none were overly annoying to me.

Here are my notes (which one of you two might want to incorporate into a reply).

1. page 2, "If one accepts this conclusion ..." : Note they take away the "a" in front of the kets. That grammatical change takes away the impression that the final state is unspecified within the set. I vote that we force them to reinstate the "a"s unless you can figure out a smoother way to express the proper idea.

2. page 2, "The physical basis of Einstein's ..." : They changed the end of the sentence to "amenable to experimental testing." That seems odd to me. Should we protest?

3. page 3, "We then use a version of the so-called Dutch-book ..." : They want us to explain "Dutch-book argument." That's pretty stupid, given that they didn't ask us to explain "Bayesian probability theory" or even "Gleason's theorem." We even used the warning sign "so-called." I don't know what more can be said without inserting Section II into the middle of this paragraph. Perhaps you guys have a nicer way to approach it than I would.

4. page 5, "The probability assignment is thus inconsistent ..." : You'll note in the proof that the equation at the end of the sentence is broken at the end of the line in an awful way. Can we ask them to keep it together? Or perhaps we can simply display the equation.

5. page 5, "For example, normalization of the probabilities ..." : This is not a problem of theirs but a question of mine. We end the sentence with "so obvious that it needs no justification." Do we really need to insert that phrase? I find it a little distracting from the main point of the sentence.

6. page 6, "The keys to these results are ..." : I would change the very last word to "paragraphs."

7. page 10, "The data gathered from the measurements are said ..." : What is this, England? "The House of Commons have voted ..." Horrible. I view "data" as a collective noun.

That's it really.

But let me take this opportunity to give one last ramble from the heart. It has nothing to do with changing the present paper: I just want to say it because it's on my mind again.

After reading the paper once more, I found myself feeling awful again. I don't think there is a reader out there besides Schack and Caves that will come away with the feeling that the Dutch-book argument is *purely* an internal consistency argument (or rationality check). The English in a sentence like,

> Given the assumptions of Gleason's theorem, if a scientist has maximal information, any state assignment that is different from the unique pure state derived in the last paragraph is inconsistent in the Dutch-book sense; i.e., it leads to a sure loss for a bet on the outcome of a measurement on a single system that includes the unique pure state among the outcomes.

is just loaded with imagery. Who out there will read "sure loss" as anything other than a factual state of affairs—something dictated by the world independent of the agent? Who out there will read the word "unique" as really meaning "unique with respect to the agent's belief of certainty"? Who out there will not interpret the phrase "maximal information" in an objectivistic way—i.e., that there is one and only one way to have maximal information? Or here's a better acid test: If a reader were to be confronted with our paper in its present form at the same time as the BFM paper, would he be able to see any philosophical differences in the approaches of the two papers? I can't imagine it. All that troubles me very deeply: A small change of language really could have made all the difference in the world.

29 November 2001, to Renes, "Observation"

Renesism 3: *here's some stuff relating to what i wrote the other day from "The Taboo of Subjectivity" by B. Alan Wallace:*

> *The disdain of scientific materialism for subjectivity has also shaped the very concept of scientific observation. While nonscientific kinds of observation also detect phenomena—such as our joys and sorrows, hopes and fears, ideas and inspirations—they are thought to be tainted by human subjectivity and are therefore suspect. From the perspective of scientific materialism, human sensory perception may be deemed not only unreliable but irrelevant. For a scientific observation to take place, all that is required is a detector, or receptor. The human eye is one type of receptor, which detects a certain range of electromagnetic frequencies, but other instruments also measure this and other types of information, and they are regarded as more reliable.*
>
> *In common parlance, for an observation to take place, the received information must be transformed into humanly accessible information that is, sooner or later, perceived and understood by a human being. But according to scientific materialism,* observation *is assimilated into the general category of* interactions*, thereby freeing it from the subjectivity of its normal associations. This interpretation is said to be central to grasping what is involved in scientific objectivity in the search for knowledge and the justification of belief.*

[reference to Dudley Shapere "The Concept of Observation in Science and Philosophy," Philosophy of Science **49** *No. 4, 1982, pg 485.]*

Beautiful quotes! Do you have the whole Shapere article? Do you know whether he is for or against that conception of observation (i.e., as interaction)? If you do have the article, could you make a copy and send it to me? And what about the Wallace thing—is that a book or an article? How is it?

29 November 2001, to Wootters, "Community"

What a beautiful letter; thank you.

In a way, I've been going to my own seminar on science and religion lately—but with me, through my reading choices. I've gotten stuck on the "pragmatism" movement, predominantly William James's version. I had never realized before what a wealth of material was there (for the kinds of thoughts I'd like to think, about quantum mechanics in particular). Nor had I realized how Wheeleresque and Woottersesque James was in his outlook—hanging so much on the idea that the universe is (in part) a product of our collective experience.

In that regard, but with respect to religion, I have never been so impressed by the possibility—necessity even—of "faith" than when I read James's articles "The Sentiment of Rationality" and "The Dilemma of Determinism." They have made a huge effect on me. If you happen to read them and have any reaction, I would love to hear your thoughts. (For fun I just had a look at the Williams College library; you guys have a collection of 63 items penned by James! In contrast, not one hit in the Bell Labs library ... but who would have guessed otherwise.)

I sympathize with your unease with telepathy. Here's how Stephen Brush put it:

> 'Wheeler's dilemma' is this: how can one maintain a strong version of the Copenhagen Interpretation, in which the observer is inextricably entangled with that which is observed, while at the same time denying that our consciousness affects that which we are conscious of—and thus accepting the possibility of telekinesis and other psychic effects? For Wheeler himself there is no dilemma at all; one simply has to recognize 'the clear distinction between (1) the strange but well verified and repeatable features of quantum mechanics and (2) the pseudo-scientific, non-repeatable and non-verified so-called extra sensory perception.' But Wheeler's own views are likely to strike a non-physicist as being just as bizarre as those of the parapsychologists he deplores. Indeed, no one has yet formulated a consistent worldview that incorporates the Copenhagen Interpretation of Quantum Mechanics while excluding what most scientists would call pseudo-sciences—astrology, parapsychology, creationism, and thousands of other cults and doctrines.

The issue whether there is something besides unitarity—whether it can ever "breakdown" as you put it—is acute, but I think a lot of the problem hinges on how one views unitarity's status in a physical theory. I am inclined to believe that unitarity—or more generally "trace-preserving complete positivity"—does not breakdown either. But that is because I am inclined to view the time-evolution mapping one ascribes to a system as an epistemological entity (rather than an ontological one), much as I view the quantum state one assigns the same system. That is to say, I am struck that there is a deep reason one can make a bijective correspondence between the completely positive maps on one Hilbert space and the density operators in a larger space, a reason that has nothing to do with imagining ancillary systems: A CP map *is* a density operator, it is a state of knowledge in just the same way any other quantum state is.

There are other (peculiarly quantum) reasons for saying what I just did above, but it is also part of a larger program I have in mind (and one I think James had in mind too)—namely, to find a little slippage between the notion of a physical theory and the world itself. Below I'll quote a little piece I wrote to Carl Caves on the idea. It gets at the sentiment, even if not at the technical details! [See 04 October 2001 note to Caves-Schack, "Replies on Pots and Kettles".]

30 November 2001, to Svozil, "Some thoughts on your paper(s)"

Thanks for your letter. You ask two good questions, and unfortunately I don't know the answer to either one! I will however have a good look at your papers and start thinking about them.

I think the Dutch-book argument for probabilities is not so strongly tied to the Boolean structure of propositions (as, say, Cox's argument for the probability axioms is). For instance, just look at Eqs. (1) and (2) in our paper. It seems to me that that part of the argument does not care one iota whether there is a distributive law for this event structure or not. If that's the case throughout the remainder of the Dutch-book mechanics, then maybe this is an important point.

05 December 2001, to Caves-Schack, "Dear Prudence"

Fellow Bayesians,

(Let's see if Carl's knowledge of 1960s music can help him guess the origin of this note's title.)

Anyway, to the real subject. Dutch-book coherence? Dutch-book consistency? Neither of them seem to be an accurate account of what's going on with the theorem. I know neither one of you will accept it, but I think "prudence" gets significantly closer to the mark.

- **Consistent** – 1. In agreement; compatible: "The testimony was consistent with the known facts." 2. Being in agreement with itself; coherent and uniform: "a consistent pattern of behavior." 3. Reliable; steady: "demonstrated a consistent ability to impress the critics." 4. In mathematics, having at least one common solution, as of two or more equations or inequalities.

- **Coherence** – 1. The quality or state of cohering, especially a logical, orderly, and aesthetically consistent relationship of parts. 2. In physics, the property of being coherent, as of waves.

- **Prudent** – 1. Wise in handling practical matters; exercising good judgment or common sense. 2. Careful in regard to one's own interests; provident. 3. Careful about one's conduct; circumspect.

- **Prudence** – 1. The state, quality, or fact of being prudent. 2. Careful management; economy.

How did Mermin's talk go? How did Ruediger's talk go?

07 December 2001, to Caves-Schack, "A Sinking Ship?"

Fellow Bayesians,

Especially in light of what Carl wrote me yesterday, perhaps you should both have a serious look at Gavriel Segre's Ph.D. Thesis at `http://xxx.lanl.gov/abs/quant-ph/0110018`. van Enk pointed it out to me this morning, and believe me it'll be worth your time! In particular, pay attention to Theorem 5.2.21 "Impossibility of a Subjectivistic Bayesian Foundation of Quantum Probability Theory" on page 193. Maybe we ought to get off this ship before it sinks?

Also, make sure you don't miss reading his acknowledgments on page 5. This is surely the best quote I've ever seen on `quant-ph`!

07 December 2001, to Caves-Schack, "Re: A Sinking Ship?"

Schackcosm 54: *Are you serious? Of course, the thought has crossed my mind.*

The sinking ship remark was a joke. The reference to the (unaccepted) thesis was a scoff (or at least a little). You did read the "acknowledgment" of the thesis, didn't you? Also look at the other contents of the "thesis."

I have never been more confident about anything scientifically than that we are on the right track in our quantum Bayesianism. No Bayesian that I've ever read called a probability assignment a "state of knowledge." What's different about quantum states? Why should they be above the "degree of belief" that every classical assignment is? Taking the "state of belief" appellation seriously is simply the cross we have to bear ... at least as I see it. (Sorry for the Christian motif, Carl.) But I'm also confident that it's not that bad anyway; it's just a question of packaging.

10 December 2001, to Caves, "Lost It"

I'll try to get in touch with you later today. Sorry we missed each other over the weekend.

Cavesism 52: *I managed to lose your last e-mail, where you expressed the view that no Bayesian views a probability as based on a state of knowledge. In the absence of your full statement, here's a thought on that point.*

I don't balk against "facts" having a say in determining probability assignments. I balk against "facts" *uniquely* determining them.

10 December 2001, to Mermin, "The Spirit of Gandhi"

Well, you took away some of my fun with the letter you just sent me! Over the weekend, I ran across your talk at the ITP and listened and watched the whole thing: In fact, I was going to send you some direct comments on it. Now some of my steam is taken away.

Still, let me gingerly point out some things I had wanted to point out earlier ... even if they may not be quite as relevant anymore.

Somewhere around 36 minutes into the talk—actually it was in your reply to a question (which sounded to come from Jeff Kimble)—you said:

Merminition 39: *Our presentation, at least in the paper we submitted on the web, can be read as being tinged with a view that quantum states are more than a reflection of knowledge ...*

the implication seeming to be that in the present talk you strove to get around that.

However, at 23 minutes into the talk you said this:

Merminition 40: *Of course there's the question of what it means for the combined knowledge of all observers to constitute a consistent body of knowledge about S, which is an interesting question. I'm taking a kind of dumb-physicist view, which is that there should be—at least in principle—there should be one observer who has a lot of data about various measurements and mutually commuting observables made on The observer having access to all the data will realize that if ...*

I know I have said all this before, but let me "focus in" on these two passages and try to say it again (per your warning about me "focusing out" too often). I would say that what you said in the second passage has nothing to do with it being a "dumb-physicist view" of things. It is, however, the sine qua non of a view "tinged" with making the quantum state more than a simple epistemic entity. For it in essence says, there SHOULD BE (your words) a "right" quantum state, or a range of "right" quantum states, and that has nothing to do with any ACTUALLY EXISTENT observers.

Please do think about the similarities between your "dumb-physicist view" and the old limerick:

> There was a young man who said, "God
> Must think it exceedingly odd
> If he finds that this tree
> Continues to be
> When there's no one about in the Quad."
>
> REPLY
>
> Dear Sir:
>
> Your astonishment's odd:
> I am always about in the Quad.
> And that's why the tree
> Will continue to be,
> Since observed by
>
> Yours faithfully, God.

I would say that in supposing THERE SHOULD BE a superobserver, you are supposing that the quantum state should—in essence—already be there without any observer at all. When Bishop Berkeley ran into trouble with the question of where the trees go when there's no observer, he invoked the idea that God was there all along. Your superobserver is in essence a God, who—through his own objectivity—re-endows the quantum state with an objectivity I THOUGHT you were trying to get rid of in the first place.

But maybe that was never your goal.

Martin Gardner said something very clear in this regard in his essay "Why I Am Not a Solipsist," and so let me quote it:

> In this book I use the term "realism" in the broad sense of a belief in the reality of something (the nature of which we leave in limbo) that is behind the phaneron, and which generates the phaneron and its weird regularities. This something is independent of human minds in the sense that it existed before there were human minds, and would exist if the human race vanished. I am not here concerned with realism as a view opposed to idealism, or realism in the Platonic sense of a view opposed to nominalism or conceptualism. As I shall use the word it is clear that even Berkeley and Royce were realists. The term of contrast is not "idealism" but "subjectivism."

(The phaneron, by the way, was C. S. Peirce's term for "the world of our experience—the totality of all we see, hear, taste, touch, feel, and smell.")

If I were to give the BFM paper and your quest to make sense of Peierls (and Bohr and Heisenberg, as evidenced by your talk) a reading, I would say that what you are trying to do is give the quantum state an idealistic interpretation (via the word "knowledge") and thinking that that somehow contrasts with a realistic interpretation (and in doing so might fix everything

up quantum interpretation-wise). But—in analogy to Gardner—for me, "realistic" and "idealistic" interpretations of the quantum state amount to the same thing. What I'm worried about is whether one can make any sense of the quantum state at all without simultaneously positing the active agent who makes use of it—I would claim you can't. And thus I'm left with the kind of "subjectivism" (for the quantum state) that frightens so many people.

PS. In case you're wondering why I titled this note "The Spirit of Gandhi," it is because I am hoping you will think of it as a form of nonviolent protest. Too many times in your talk you pointed out what a violent reaction I had to the BFM result, and I just can't think that that helped my reputation as a rational thinker or the seriousness with which a Bayesian-kind of quantum-foundation attempt should be viewed. (But I love you just as much as ever—maybe more with all the advertising you gave me—and there are no grudges.)

10 December 2001, to Mermin, "Trumps and Triumphs"

Also,

Merminition 41 : *I did urge Carl and Ruediger to do something about the phrase (in your other paper) "Gleason's theorem is the greatest triumph of Bayesian reasoning" which I read as a claim that Bayesian reasoning is the only way (or at least far and away the best way) to derive and understand Gleason's theorem. They explained that what you really meant was that Gleason's theorem, by providing stringent constraints on possible prior distributions, provided a powerful tool for Bayesian reasoning. But I think they did an inadequate job of removing the ambiguity, which is too bad, since I am rather sure it will put off most of the people you ought to be addressing.*

I'm glad you urged them. Here's what I had said to Carl and Ruediger in a September 20 note:

> Another thing I guess I really didn't like—and this is only a new one for me, it's not something that had eaten away at me before—is the slogan "Gleason's theorem can be regarded as the greatest triumph of Bayesian reasoning." To say that is to imply that Bayesian reasoning IMPLIES Gleason's theorem. I don't think you mean that. I think what you mean is that it is one of the most valuable additions to Bayesian reasoning ever. But I won't cause trouble here: I know you both like the saying.

But they generally just view me as a trouble-maker.

10 December 2001, to Mermin, "Reread (pronounced re-red)"

I'm just reading through all my emails of the day again.

Merminition 42 : *P.S. Spent a lot of time talking with Carl and Ruediger in SB, which I enjoyed very much. I think the first half of your draft joint paper formulates the basis for the "necessary" condition of BFM much more coherently than we did: i.e. it is necessary if there is to be any density matrix that does not contradict anybody's strongly held beliefs. Much better than talking about "subsets of a body of data that anybody with access to all of which would agree constitutes nothing but valid results of measurements, none of which invalidate earlier measurement outcomes...." It's hard even to put into a grammatical sentence. The other four cases strike me as interesting extensions of the case we (and Peierls) were addressing, which concerned existing knowledge prior to any subsequent interventions. My feeling was that there was no serious disagreement among the three of us, but maybe they will give you different reports.*

Thanks, that's a very nice compliment. Here's another way to put it, which might be my preferred way. Suppose two strongly consistent agents want to stand a chance of coming to agreement, just by talking to each other. Then BFM is the necessary condition. If it is not satisfied, then the only way they will ever get anywhere is by "consulting" the system with a (mutually agreed upon) measurement: With it, they can bend the world into something more congenial to both.

14 December 2001, to Caves-Schack, "Strong Consistency"

I know that I'm terribly slow to be coming up to speed on this project. But now I have another possibly frivolous question.

What was the origin of strong consistency? What motivated Ruediger to invent the concept? And indeed, did either of you ever find it appearing in the literature otherwise? It, of course, has a very different flavor than normal Dutch-book—I know that's no surprise to you—but for me, for the moment anyway, it seems to me to be introduced more for mathematical nicety than anything else. (I.e., It seems more to be introduced just for the purpose of cleaning things up a bit axiom-wise.) It doesn't have the same operational/pragmatic flair that the rest of Dutch book has.

Indeed, it might be useful to lay this sort of issue on the line in the paper.

I can think of at least one case where one would never want to enforce strong consistency, namely for infinite event spaces. Consider an infinite number of draws from an i.i.d. distribution. One has that the probability of a typical sequence is one; however one should never then say that that implies the bettor believes a typical sequence will occur with certainty.

Reading ever so slowly (but I hope thoroughly),

15 December 2001, to Caves, "Bah, Humbug"

Humbug. Did I ever tell you I hate Christmas: It's the most intrusive holiday I know of, with people pushing me all over with obligations that they've invented for me. (Susie, Kiki's mom, arrived yesterday, by the way.) Now Thanksgiving, that's my favorite holiday: you cook, you eat, you watch a parade, and you end up rubbing your belly in a deep soporific ecstasy.

By the way, speaking of rubbing bellies, Kiki went into false labor last night. At first, when it was over, I thought, "Thank God it's false." But now, in the light of the morning, I'm thinking, "When are we ever going to get this over with?"

Cavesism 53: *There is evidence that it's out there in the literature (Sklar mentions it without any references), but we haven't dug up any references. In fact, Ruediger and I had the impression that you might know where to find references.*

This is one of my biggest gripes about not being in a university. I don't have easy access to anything outside of the physics and mathematics journals. (In fact strangely, last night I dreamed that I was back at the University of Texas, surrounded by books, books, books on all subjects, and I was telling myself how glad I was to be back.)

Did you have a look at the Kemeny and Lehman references I recommended in an October 27 note? That at least would be my starting point. Also Bernardo and Smith, as I recall, had a lot of references on Dutch-book—in particular a lot of references dissatisfied with the argument (as Bernardo and Smith, in fact, are).

Cavesism 54: *I don't know about the flair point. I'll grant that strong consistency isn't as compelling as ordinary consistency, but it still registers on the flair scale: in a single trial, you never*

win, but definitely sometimes lose. Moreover—and this is the crucial point—strong consistency is absolutely essential so that probability assignments and density operators incorporate sufficient information about one's beliefs that one can use them as surrogates for the beliefs in our arguments.

I wouldn't say that it's not "as compelling", it's just of a different flavor. To say that it's not as compelling, means it's directly comparable to the Dutch-book argument. But it's not clear to me that it is. Presently it seems to me much more like a Cox kind of thing. It says (at least in one case) my outward commitment *must* reflect my internal belief—not because I will be imprudent if I don't do so, but because I don't want my friends to see any discrepancy between my beliefs and my actions.

I believe I understand what you call the "crucial point." But I'm guessing—just a psychological point—you're more attracted to it, because it gives rise to (a better version of) the BFM statement. Thus it gives the impression of giving something very solid. Whereas simple Dutch book lets everything fly in the wind—maybe that's the very thing I'm attracted to it.

In any case, I don't see anything wrong with exploring what strong consistency leads to. I'm just trying to get it straight in my head—and possibly also for the reader—what it's all about.

Cavesism 55: *Chris said, "Indeed, it might be useful to lay this sort of issue on the line in the paper."*

Perhaps you could be more specific. It seems to me that it is pretty well laid out in the paper, but you're looking at it with a fresh eye, so suggestions are welcome.

I'm sorry, I don't quite know how to be more specific right now. Maybe I would just like the reader to see more of the debate about it. And it should come across that it is (possibly) an invention of C and S, not de Finetti or Ramsey.

Cavesism 56: *You're absolutely right here: probability 1 cannot mean certainty and probability 0 cannot mean impossibility in the limit of an infinite number of draws from an i.i.d., but it's not clear to me that the problem is strong consistency. To apply probabilities to infinite sets, you need countable additivity for exclusive events, which doesn't follow from the finite additivity given by the Dutch book. To apply probabilities to continuous sets requires measure theory. These things are additional mathematical structure beyond what the Dutch book can provide. Perhaps the best philosophy is Ruediger's: probabilities really apply only to finite sets of events; the generalization to infinite sets is an idealization, where additional mathematical structure is added to make things work nicely.*

Point well taken. But on the other hand, I'm still not convinced either. I think all three of us are aligned that the only thing really worth considering, conceptual-wise that is, are finite event spaces. However, it still strikes me as a chink in the armor that I can think of a limiting case where I would not want probability $1 - \epsilon$ to have anything to do with internal certainty. That is to say, maybe there's a deep reason that Dutch-book consistency is NOT strong consistency ... and maybe we're missing something by plastering over its unsightliness so quickly.

Cavesism 57: *Let me know if you have specific recommendations. Ruediger and I have agreed on a number of changes in presentation and a large number of additions on the POVM front. I will be attempting to incorporate these this weekend.*

I'm trying my best in the light of the present events. (But I'm sure you're not holding your breath for me.) The main thing I'm shooting for now, is to make sure I understand every aspect of the paper—in a defendable way, that is—before Mermin arrives Tuesday.

Did, by the way, you understand my issues with the *sufficiency* of the det = 0 conditions? Was I missing something?

19 December 2001, to the world, "Katie Viola Fuchs"

Dear Family, Friends, and Colleagues,

Katherine Viola Fuchs came to the world with Monday's sunrise, 17 December 2001. She was 8 pounds, 19.5 inches, and a delight to her parents' eyes—she is beautiful in every way. And like her big sister Emma, Katie has already made it clear that the future will not just unfold before her, but be made in crucial part by her presence and will.

With the greatest of expectations (and a modicum of pride), we send greetings to everyone on Katie's behalf.

Chris, Kiki, and Emma Fuchs

15 December 2001, to Caves, "One So Far"

Just looking at the end of Section III, I was reminded that BFM posted there final version of the paper on `quant-ph` this week. So, they have made some changes. It might be worthwhile to check that they didn't sneak in too drastically different a formulation of the question.

Also, there you (or Ruediger) write, "In contrast, our derivation is couched wholly in terms of the beliefs of the parties and does not appeal to a real state of affairs; it is therefore preferable in a Bayesian approach to quantum mechanics." I'm not sure—at least as hinted by the present structure of the manuscript—that any nonBayesian we have a clue why it's a preferable statement of the problem. "What's wrong with letting the quantum state reflect an objective fact, that what cannot happen cannot happen?," the reader might ask. "Aren't these CFS guys just paining over inconsequential points?"

Personally, I would like to see some discussion, perhaps in the introduction or in Section III, along the lines of those Bernardo and Smith quotes that I like so much. I'll place them below in case you forgot. I think they give the whole motivation for the work.

Finally, the present manuscript gives pretty short shrift to the Peierlsian ideas that got David's blood up in the first place. Nor do we say much about the apparent inconsistency between the goals and the "solution" of BFM in relation to the Peierls quest.

02 January 2002, to Bennett-Smolin, "Solipsism Story"

Here's the other story from Martin Gardner's "Why I Am Not a Solipsist."

> Russell once spoke on solipsism at a meeting chaired by Whitehead. As Russell tells it in his autobiography, he said he could not believe he had written those parts of Whitehead's books which he (Russell) could not understand, although he could find no way to prove he hadn't.

04 January 2002, to Caves, "New Breach of Faith"

Cavesism 58: *I agree wholly with your statement in the following:*

Also, there you (or Ruediger) write, "In contrast, our derivation is couched wholly in terms of the beliefs of the parties and does not appeal to a real state of affairs; it is therefore preferable in a Bayesian approach to quantum mechanics." I'm not sure—at least as hinted by the present structure of the manuscript—that any nonBayesian [would] have a clue why it's a preferable statement of the problem. "What's wrong with letting the quantum state reflect an objective fact, that what cannot happen cannot happen?," the reader might ask. "Aren't these CFS guys just paining over inconsequential points?"

But I don't quite know what to do about it. Your B& S quotes are nice, but they will just get in the way, in my view, of getting people even to grasp the setting for what we're doing.

Perhaps you misunderstood: I am not asking that the Bernardo and Smith quotes actually be used in the text. What I am asking is that we find a way to convey to the reader why any of this is important. One thing you've got to realize is that more than once, I've heard Mermin describe the reception of the BFM paper with words like these:

Merminition 43 : *Sorry, I had the feeling that those few who understood anything at all thought it was pretty obvious, but they were polite anyway.*

What—it seems to me—is our task, is to convey to any readers who had such a reaction as above that they were actually snookered by BFM. The Bernardo and Smith report of what Bayesianism is about builds a context for our efforts—they make it clear why one should expect a hierarchy of conditions (like the ones we explore in the paper) rather than an *absolute* answer.

It is in saying terse, "only established clique"-interpretable (i.e., only preformed-radical-Bayesian interpretable) things like, "In contrast, our derivation is couched wholly in terms of the beliefs of the parties and does not appeal to a real state of affairs; it is therefore preferable in a Bayesian approach to quantum mechanics" ... and just kind of leaving it at that, without any further buffer ... that is going to get us in trouble. Or at least, that is one of my continuing fears whenever I work with you and Ruediger.

What is being laid in this paper is the groundwork for viewing the quantum state in a way that people—even the Bayesians among us—are not at all accustomed to: Namely, taking the quantum state's subjectivity absolutely seriously and to the extreme. I can't understand how leaning heavily on the motivation for this work can hurt the paper. Indeed it seems far more crucial than the technical results in Section V if you ask me: if no one cares about the results, no one will read them in the first place. (It's not like this stuff can, with almost a single word, be advertised as a new quantum algorithm like Shor's.)

What's a little annoying is that I know that you (from among the three of us) are the one most up to the task of writing a beautiful, yet businesslike introduction. The only thing it seems that you need to be convinced of is that people will actually read this paper if it is well-written (and won't read it otherwise). Strangely, enough, my wacky 45-page papers get read (or at least skimmed): Explain that. It's not the substance a priori (and maybe not even a posteriori), but it might be the style.

Finally, in a last-ditch effort to shore up this point, I'm going to perform a new breach of faith (and just reconcile it with St. Peter when the day comes). [...]

OK, I'll say no more on the subject.

Cavesism 59 : *Example: Kimble sat through Ruediger's talk at Caltech, nodding his head in agreement the whole time, but then it emerged in the discussion afterward that he thought each party had his own copy of the system. He and his student, Andrew Landahl, really had a hard time*

with understanding that there is any [difference] between "different states" and "different systems." They really just couldn't help using these interchangeability.

This is an oddity I have encountered before with several people. I don't understand its origin, but it is weird. Even taking an extreme realist view that there are objective quantum systems and objective quantum states, one should be able to detect the matter-property dichotomy in that and not confuse the two. If you have a theory of the confusion's origin, I'd be interested to hear it.

04 January 2002, to Caves-Schack, "Once Again"

Cavesism 60: *It IS interesting that strong consistency is required so that the firm part of your belief structure can be read off your probability assignments. I would have thought that a Bayesian would want to know what assumptions are required to translate beliefs into probability assignments.*

Let me put the relevant part of my take on that down again:

> It [strong consistency] says (at least in one case) my outward commitment *must* reflect my internal belief—not because I will be imprudent if I don't do so, but because I don't want my friends to see any discrepancy between my beliefs and my actions.

"Outward commitment" means acceptable odds for the bettor. What strong consistency says is that there are cases I should lay my precise beliefs on the table (for public view), even if I didn't have to in order to avoid a sure loss.

Nowhere else does Dutch book do this. For instance—with standard Dutch book—I could internally believe $p(X, Y)$, but nevertheless only accept bets according to $q(X, Y)$ so long as both these are coherent assignments. Is it part of the Bayesian creed to also require honest reporting of my internal beliefs? Maybe it is: but it seems to me that that is something on top of Dutch book (and in fact we know that it is).

I'm not "against" strong consistency, I just want to understand what motivates it other than that it leads to more equations in our paper ... and thus leads to a more scientific look.

I'm going to try to get to a real library this weekend and dig up those J. Symb. Logic papers by Shimony etc. Maybe that'll demystify things for me.

07 January 2002, to Mermin, "Correlation without Correlata"

Let's see what kind of reaction the longer quote below gets out of you. "There are, so to speak, relations all the way down, all the way up, and all the way out in every direction: you never reach something which is not just one more nexus of relations."

> In the rest of this essay I shall be trying to sketch how things look when described in antiessentialist terms. I hope to show that such terms are more useful than terminologies which presuppose what Dewey called 'the whole brood and nest of dualisms' which we inherit from the Greeks. The panrelationalism I advocate is summed up in the suggestion that we think of everything as if it were a *number*.
>
> The nice thing about numbers, from my point of view, is simply that it is very hard to think of them as having intrinsic natures, as having an essential core surrounded by a penumbra of accidental relationships. Numbers are an admirable example of something which it is difficult to describe in essentialist language.

To see my point, ask what the essence of the number 17 is—what it is *in itself*, apart from its relationships to other numbers. What is wanted is a description of 17 which is different *in kind* from the following descriptions: less than 22, more than 8, the sum of 6 and 11, the square root of 289, the square of 4.123105, the difference between 1,678,922 and 1,678,905. The tiresome thing about all *these* descriptions is that none of them seems to get closer to the number 17 than do any of the others. Equally tiresomely, there are obviously an infinite number of other descriptions which you could offer of i7, all of which would be equally 'accidental' and 'extrinsic'. None of these descriptions seems to give you a clue to the intrinsic seventeenness of 17—the unique feature which makes it the very number that it is. For your choice among these descriptions is obviously a matter of what purpose you have in mind—the particular situation which caused you to think of the number 17 in the first place.

If we want to be essentialist about the number 17, we have to say, in philosophical jargon, that *all* its infinitely many different relations to infinitely many other numbers are *internal* relations—that is, that none of these relations could be different without the number 17 being different. So there seems to be no way to define the essence of seventeenhood short of finding some mechanism for generating *all* the true descriptions of 17, specifying all its relations to *all* the other numbers. Mathematicians can in fact produce such a mechanism by axiomatizing arithmetic, or by reducing numbers to sets and axiomatizing set theory. But if the mathematician then points to his neat little batch of axioms and says, 'Behold the essence of 17!' we feel gypped. There is nothing very seventeenish about those axioms, for they are equally the essence of 1, or 2, of 289, and of 1,678,922.

I conclude that, whatever sorts of things may have intrinsic natures, numbers do not—that it simply does not pay to be an essentialist about numbers. We antiessentialists would like to convince you that it also does not pay to be essentialist about tables, stars, electrons, human beings, academic disciplines, social institutions, or anything else. We suggest that you think of all such objects as resembling numbers in the following respect: there is nothing to be known about them except an initially large, and forever expandable, web of relations to other objects. Everything that can serve as the term of a relation can be dissolved into another set of relations, and so on for ever. There are, so to speak, relations all the way down, all the way up, and all the way out in every direction: you never reach something which is not just one more nexus of relations. The system of natural numbers is a good model of the universe because in that system it is obvious, and obviously harmless, that there are no terms of relations which are not simply clusters of further relations.

To say that relations go all the way down is a corollary of psychological nominalism: of the doctrine that there is nothing to be known about anything save what is stated in sentences describing it. For every sentence about an object is an explicit or implicit description of its relation to one or more other objects. So if there is no knowledge by acquaintance, no knowledge which does not take the form of a sentential attitude, then there is nothing to be known about anything save its relations to other things. To insist that there is a difference between a nonrelational *ordo essendi* and a relational *ordo cognoscendi* is, inevitably, to recreate the Kantian Thing-in-Itself. To make that move is to substitute a nostalgia for immediacy, and a longing for a salvatory relation to a nonhuman power, for the utopian hope which pragmatism recommends. It is to reinvent what Heidegger called 'the ontotheological tradition'.

For psychological nominalists, no description of an object is more a description of the

'real', as opposed to the 'apparent', object than any other, nor are any of them descriptions of, so to speak, the object's relation to itself—of its identity with its own essence. Some of them are, to be sure, better descriptions than others. But this betterness is a matter of being more useful tools—tools which accomplish some human purpose better than do competing descriptions. All these purposes are, from a philosophical as opposed to a practical point of view, on a par. There is no over-riding purpose called 'discovering the truth' which takes precedence. As I have said before, pragmatists do not think that truth is the aim of inquiry. The aim of inquiry is utility, and there are as many different useful tools as there are purposes to be served.

Common sense—or at least Western common sense—has trouble with the claim that numbers are good models for objects in general because it seems counterintuitive to say that physical, spatiotemporal objects dissolve into webs of relations in the way that numbers do. When numbers are analysed away into relations to other numbers, nobody mourns the loss of their substantial, independent, autonomous reality. But things are different with tables and stars and electrons. Here common sense is inclined to stick in its toes and say that you cannot have relations without things to be related. If there were not a hard, substantial autonomous table to stand in relation to, e.g., you and me and the chair, or to be constituted out of hard, substantial, elementary particles, there would be nothing to get related and so no relations. There is, common sense insists, a difference between relations and the things that get related, and philosophy cannot break that distinction down.

The antiessentialist reply to this bit of common sense ...

08 January 2002, to Mermin, "Re: Correlation without Correlata"

Merminition 44: *I don't like the stuff about 17 — very unconvincing. But after that it gets more interesting. It starts to get most interesting just as you cut it off.*

Of course it's not convincing: philosophical mumbles—it seems to me—can't serve that purpose. I was just curious whether it struck any chords with you on how to convey (whatever it is you've been trying to get at with) your "correlation without correlata." Is it a good metaphor? Does it carry any Ithacan soul?

Personally, I thought the analogy was nice from the second I read it. Even the primeness of a number, for instance, can only be defined by invoking the existence of all the other numbers. Numbers just don't have any properties in and of themselves (or at least none that I could think of).

You're not going to QIP at IBM next week, are you?

08 January 2002, to Mermin, "Information → Belief → Hope ??"

I started reading through the essay again and realized that if I were to fulfill your request, I might just have to copy the whole damned article. I give up! But at least I was kind enough to check that you have the source in your library at Cornell (call number and location below).

The essay is titled "A World without Substance or Essences." It should probably be read in conjunction with the essay previous to it, "Truth without Correspondence to Reality." The book as a whole is entertaining, but it is a little skimpy on firm argument. (Though, the author admits

it is an attempt at popularization ... so that kind of makes it OK in my mind.) He's certainly not the devil that Steve Weinberg labeled him. But I will admit, his version of pragmatism may go too far for my tastes: I'll hold on to the final verdict for a while. (William James is still the best bet in my eyes.)

PS. The title to this note is a joke, based on Rorty's own title (and some of the discussions in his book). However, I am seriously toying with the idea of making a distinction between "beliefs" and "commitments." I.e., saying that a quantum state ascription is explicitly a "commitment" rather than a "belief"—a commitment to behaving one way or the other in the face of some experimental data (yet to be gathered). Sometimes commitments explicitly correspond to beliefs (as Dutch-book takes to be a definition), but it seems to me not always the case.

But as I say, I'm just toying with the idea: I know Caves and Schack will have a cow and beat me up if I have enough nerve to say anything about it. So, I'd better be sure of myself.

08 January 2002, to Caves-Schack, "Re: Help with ET Quote"

Cavesism 61: *I've been trying to remember without success where ET discusses his ability to toss coins a la what we saw from Dan Greenberger in Vaxjo. Can you help?*

I don't know that I've ever actually read the example: I've just heard stories of it (probably through Ruediger), and maybe words to the affect as a section heading or something. I have a faint memory that it was somewhere in Jaynes's big book, but I couldn't it in the table of contents. Ruediger's probably a better person to ask.

One interesting thing I did find in looking for the answer to your question though, is a Jaynesian diatribe against a Dutch-book foundation for probability theory. It's in Appendix A of his book. The fear he expresses there strikes me as not so different from what I was reading into your worries about my fall into "radical Bayesianism."

08 January 2002, to Caves-Schack, "Term Origin"

Here's a little (unfortunately inconclusive) discussion on the origin of the "Dutch book" term: `http://www.fee.uva.nl/creed/wakker/miscella/Dutchbk.htm`.

08 January 2002, to Nicholson, "Rorty on Religion"

I wish I had had these quotes in stock when we were having our conversation the other night. But I hadn't gotten that far in the book yet ... if I were only a quick reader like you. [Disclaimer: My copying these quotes for your thought (and our continued discussion) neither represents an endorsement for or against their content.]

From the essay "Religious Faith, Intellectual Responsibility and Romance" in the R. Rorty, *Philosophy and Social Hope*:

page 153:

> If one accepts that claim, one will have reason to be as dubious as James was of the purportedly necessary antagonism between science and religion. For, as I said earlier, these two areas of culture fulfill two different sets of desires. Science enables us to predict and control, whereas religion offers us a larger hope, and thereby something

to live for. To ask, 'Which of their two accounts of the universe is true?' may be as pointless as asking, 'Is the carpenter's or the particle physicist's account of tables the true one?' For neither question needs to be answered if we can figure out a strategy for keeping the two accounts out of each other's way.

page 156:

Pragmatists are not instrumentalists, in the sense of people who believe that quarks are 'mere heuristic fictions'. They think that quarks are as real as tables, but that quark talk and table talk need not get in each other's way, since they need not compete for the role of What is There Anyway, apart from human needs and interests. Similarly, pragmatist theists are not anthropocentrists, in the sense of believing that God is a 'mere posit'. They believe that God is as real as sense impressions, tables, quarks and human rights. But, they add, stories about our relations to God do not necessarily run athwart the stories of our relations to these other things.

Pragmatist theists, however, do have to get along without personal immortality, providential intervention, the efficacy of sacraments, the Virgin Birth, the Risen Christ, the Covenant with Abraham, the authority of the Koran, and a lot of other things which many theists are loath to do without. Or, if they want them, they will have to interpret them 'symbolically' in a way which MacIntyre will regard as disingenuous, for they must prevent them from providing premises for practical reasoning. But demythologizing is, pragmatist theists think, a small price to pay for insulating these doctrines from 'scientific' criticism. Demythologizing amounts to saying that, whatever theism is good for, it is not a device for predicting or controlling our environment.

page 157-158:

I said earlier that many readers of 'The Will to Believe' feel let down when they discover that the only sort of religion James has been discussing is something as wimpy as the belief that 'perfection is eternal'. They have a point. For when Clifford raged against the intellectual irresponsibility of the thesis, what he really had in mind was the moral irresponsibility of fundamentalists — the people who burnt people at the stake, forbade divorce and dancing, and found various other ways of making their neighbours miserable for the greater glory of God. Once 'the religious hypothesis' is disengaged from the opportunity to inflict humiliation and pain on people who do not profess the correct creed, it loses interest for many people. It loses interest for many more once it is disengaged from the promise that we shall see our loved ones after death. Similarly, once science is disengaged from the claim to know reality as it is in itself it loses its appeal for the sort of person who sees pragmatism as a frivolous, or treasonous, dereliction of our duty to Truth.

09 January 2002, to Mabuchi, "Doin' the Dutch-Book Zombie"

Hey MacArthur boy,

It's a funny thing being a parent of a newborn child. It leads to a kind of zombie-like state most every night: You're never neither really awake nor asleep. And then your mind gets hung up on some little thing—tonight being the Dutch-book argument—and you just repeat it over and over, as if a trance. Does that build any imagery for you?

Anyway, in looking up Andrew's email address last night (or was it this night? — time stops), I ran across your web page. The flattery to be listed among your collaborators! (Keep it there!) But you need to get the affiliation right — Bell Labs, Lucent Technologies.

Philosophically, lately, I've been taken away by William James and John Dewey. I've been reading their stuff with a pretty voracious appetite (and there is a heck of a lot of it). But of course one thing leads to another—just like that marijuana—and last week I found myself picking up a copy of Richard Rorty's *Philosophy and Social Hope*. His flavor of pragmatism goes maybe too far even for me, but he writes well and I find it easy to read him. I guess I just write you all this to tell you he speaks highly of your friend Derrida! One of these days, I am going to get the nerve to approach that man. (By the way, Mermin told me recently that your old professor Bas vF has been reading my papers ... and, apparently, disagreeing with them ...)

I'm going to slip back under the covers now; like Nosferatu I keep a little fresh earth there for comfort.

Just a sleepy note to let you know I miss you sometimes.

He did the mash
He did the monster mash
The monster mash
It was a graveyard smash

13 January 2002, to Mabuchi, "Princeton Envy"

[Mabuchi said,] "sounds like domestic life is treatin' you good!"

Wasn't it Freud who said, once a school girl sees your Princeton, she'll know what's missing from her life and be envious thereafter?

Well I'm not a school girl, but I saw your Princeton the other day (ahem), and it certainly did start a deep yearning in me. What a wonderful place! This was my first real trip there since moving to New Jersey, and I was enchanted all day. I found myself thinking, what I wouldn't give to live the rest of my life in this little cloister.

The thing that really struck me was the immense resources at one's finger tips. I found myself copying a little over $35 of articles in Firestone Library! I couldn't believe it: They had the complete collection of the *Danish Yearbook of Philosophy* and (shockingly) 35 years of the *Transactions of the Charles Sanders Peirce Society*, and I knew that those were just the tip of an iceberg. I was in heaven.

Emma and I play this game: I say, "When you go to college, I hope you'll go to Harvard." She says, "I want to go to Princeton." Or if I say Princeton, she says Harvard—it's always the opposite. (You can see a trend in our relationship starting to form.)

But that's just an aside (to tell you that the grass is always greener). Keep up the good work with all those good students. Get them to read William James' *Pragmatism* and tell them that quantum mechanics is a much better motivation for all that he said there ... but to never lose sight that the real goal is to get to where he wanted to go.

13 January 2002, to Mermin, "Reality in the Differential"

Merminition 45 : *Actually I'm leaving it in the original form for now, but You're giving me a very hard time here.*

Let me say touché, and then little more than that. Strangely, actually, I've been having a little conversation of my own with You this week. And I've been planning to write you—i.e., the one with the little y—all about it. But what happens? Now that I've got a little time this Sunday morning, I'm finding that the inclination is leaving me.

So this note is going to come out far weaker than I had planned. The main thing was to build a conversation around another Rorty quote, and to tell you how pleased I am with the pragmatism movement in general. I'm finally finding a philosophy so close to what I'm looking for that I'm willing to advertise it to my friends.

Somehow they go a little too far for me, but there are so many beautiful gems I keep finding in their stuff that I find it better not to dismiss it outright. The second paragraph below struck me especially last week. It might as well be about my latest quantum interpretation thoughts.

The reality is in the differential. The quantum state represents a (gambling) commitment on the part of the agent; it never represents anything external to that agent. If you're looking for where the "reality" of the external world creeps into the formalism, you should look to how these commitments change. That's what I've been trying to say for a few months now ... but I guess you already understood that (as maybe witnessed by our last real conversation, during your visit to Morristown).

I had a pretty happy-sad week this week you might say. The happiness was that I visited Princeton for the first time since moving up here. I was like a kid in a candy store!! I ended up copying over $35 of articles in their wonderful library (which even subscribes to the *Danish Yearbook of Philosophy*!). The sadness was in that I had to leave that environment and go home at the end of the day; it wasn't home itself.

I'm going to write you that longer James-Dewey note eventually. But right now I'll just leave you with a reminder that the reality is in the differential. From: R. Rorty, "The Pragmatist's Progress: Umberto Eco on Interpretation," in his book *Philosophy and Social Hope*, (Penguin Books, New York, 1999), pp. 131–147:

> As I see it, the rocks and the quarks are just more grist for the hermeneutic process of making objects by talking about them. Granted, one of the things we say when we talk about rocks and quarks is that they antedate us, but we often say that about marks on paper as well. So 'making' is not the right word either for rocks or for marks, any more than is 'finding'. We don't exactly make them, nor do we exactly find them. What we do is to react to stimuli by emitting sentences containing marks and noises such as 'rock', 'quark', 'mark', 'noise', 'sentence', 'text', 'metaphor' and so on.
>
> We then infer other sentences from these, and others from those, and soon—building up a potentially infinite labyrinthine encyclopedia of assertions. These assertions are always at the mercy of being changed by fresh stimuli, but they are never capable of being *checked against* those stimuli, much less against the internal coherence of something outside the encyclopedia. The encyclopedia can get *changed* by things outside itself, but it can only be *checked* by having bits of itself compared with other bits. You cannot *check* a sentence against an object, although an object can *cause* you to stop asserting a sentence. You can only check a sentence against other sentences, sentences to which it is connected by various labyrinthine inferential relationships.
>
> This refusal to draw a philosophically interesting line between nature and culture, language and fact, the universe of semiosis and some other universe, is where you wind up when, with Dewey and Davidson, you stop thinking of knowledge as accurate representation, of getting the signs lined up in the right relations to the non-signs. For you also stop thinking that you can separate the object from what you say about it,

the signified from the sign, or the language from the metalanguage, except *ad hoc*, in aid of some particular purpose.

21 January 2002, to Peres, "R, B, and P"

Asherism 2: *I am happy that QIP2002 went well. I saw the papers of Briegel and Nielsen on quant-ph (please remind me the numbers, if you have them ready). I was not favorably impressed, maybe I misunderstood them, and I should read them again.*

Actually, the better work is the Raussendorf/Briegel stuff. If you were not favorably impressed, then I think you should give it another chance. I think it is a beautiful construction, and, in fact, the deepest thing I've seen in quantum computing for 2 or 3 years now. The first paper to start with is `quant-ph/0010033`. Then more details can be found in `quant-ph/0108067`, `quant-ph/0004051`, and `quant-ph/0108118`.

What is deep about this work is that all computations start off with the SAME given entangled state for the qubits. That is to say, a given entangled state is taken as a resource for the task. Thereafter the particular computation one is interested in is enacted by making single-qubit measurements alone: there are no further unitary evolutions. I think that is quite remarkable and quite lovely.

An interesting feature of the Raussendorf/Briegel model (as opposed to the Nielsen and Leung models) is that one need to take into account NO details of the post-measurement state for the measured qubits: they can be thrown away immediately after the measurement. And because the measurements are localized, the post-measurement state of the remainder of the qubits is fixed completely by the POVM (rather than the operation). That is to say, for the relevant qubits in this model, "an effect only has one operation."

Also: Again, I apologize about the typo that caused concern over Petra.

27 January 2002, to Nicholson, "Re: What Would William James Say? 1."

Of course I'm reading email in Texas! (Or at least that's what I'd say.) I'll send your regards to my Mom.

El Jeffy 1: *"the universe we observe has precisely the properties we should expect if there is, at bottom, no design, no purpose, no evil and no good, nothing but blind, pitiless indifference."*

That may have been true at one time, i.e., at some stage in the development of the world. That's how it got off to a start so to speak. But now, I go the experimentalist (who does the hard work in helping the theorist construct his theories), and ask him, "What ya doin?" He says, "Twiddlin' knobs." I ask, "How come?" He says, "I'm tryin' to hep this wacky friend of mine. He wants to get a theory of how iodine reacts to this and that? I'm chartin' it out, givin' him some clues." I ask, "How do you do that?" He says, "By twiddlin' these knobs."

The present stage of the universe doesn't look like it's full of blind indifference to me.

The initial condition is always left separate from the theory; there's a reason for that. How else would the experimentalist be able to twiddle his knobs?

27 January 2002, to Nicholson, "Re: What Would William James Say? 2."

El Jeffy 2: *Anyway, I was far more interested in your reaction the quote requiring God's absence from standard scientific theories.*

The point was, it doesn't even require the scientific agent's absence from the ultimate gears and pinions of the world, much less a god's. The scientific world view consists essentially of two components: theories and initial conditions. And, it seems to me, it is a tacit assumption of the whole scientific enterprise that the experimentalist can freely set the initial conditions he wishes to.

El Jeffy 3: *If the universe really wasn't blindly indifferent, people wouldn't fly airplanes into buildings, convinced they were on their way to heaven.*

There's something about this sentence I just don't like. I'll try to put my finger on it in the next couple of days.

29 January 2002, to Renes, "Qunix"

Renesism 4: *in other news, i discovered a really interesting paper (actually an undergrad thesis) by a guy at oxford, who argues against deutsch's view of the quantum world. one thing he takes to task is the notion that information is physical (which i think i've now discarded).*

Yeah, I met Timpson in Ireland and really enjoyed his company. I'll definitely look at his thesis. Right before Rolf died, I had wanted to write him a letter telling how much I had come to disliking the phrase "information is physical." I wanted to tell him that I think a far more appropriate phrase would be "information carriers are physical." In fact, I told Charlie Bennett about this—at the time—and he told me, "Too late; Rolf just had some fraction of his brain removed last week."

Concerning myself, I've gotten further carried down the path of pragmatism. I've even read some of the Rorty blend now. Indeed my latest little epiphany hit me last week (during an operating-systems talk in our center) when I came up with the following slogan: "A physical theory really amounts to little more than a programming language." Its rules, its specifications are more about the ways we've come to naturally manipulate the world than anything intrinsic to that same world. I tried to say this in my Oct 4 letter to Carl ("Replies on Pots and Kettles" in the new mini-samizdat), but I think the new slogan says it better.

29 January 2002, to Caves, "Re: Växjö Contribution"

You know, if my true love is philosophy—Bennett calls it theology in my case, actually—yours is certainly sports. You're just not going to let this go without one hell of a fight, are you?

Cavesism 62: *So it appears that I'm in a pickle with evolutions that are mixtures of unitaries, since quantum operations don't have unique decompositions. But of course, all the decompositions into things other than unitaries aren't of interest, and I've been able to "show" (a number of half-baked steps here) that mixtures of unitary EVOLUTIONS are unique, at least in a sense that's good enough for me. To put it more precisely, I've shown it for qubits and think I can go further.*

I'm wondering what the "sense that's good enough for me" is? For instance, I already think of the depolarizing channel, where it doesn't matter which x, y, and z axes I use for defining the Kraus operators. But you probably have something up your sleeve that will be more instructive than that.

Cavesism 63: *For pure states the problem shows up as an inability to make a clean distinction between objective and subjective probabilities,*

And for the channel maps, I would say the problem shows up as an inability to make a clean distinction between objective and subjective probabilities. [Just like my double footnote in the NATO paper, I meant this to be taken seriously.] The only difference now is that the probabilities are of a conditional type, $p(y|x)$. But I know that this is too cryptic for you to make any sense of it at the moment, and I already write you too much preachy email: I just need to try to write a damned paper and be done with it.

02 February 2002, to Timpson, "Colleague"

Well, I still haven't read your anti-Brukner/Zeilinger paper, but I'm writing this small note to tell you that I just finished reading your undergraduate thesis *Information and the Turing Principle*. It's quite a work, and I very, very much enjoyed reading it!

I have to tell you, one of my first reactions after reading about the first third of the thesis was, "Finally there's something sensible coming out of Oxford!" It really was such a relief: I had been thinking that Deutsch had essentially brainwashed everyone in the quantum information community there (except for possibly Hardy and Steane).

Of course, what I like most is what I see as a significant overlap between our attitudes toward scientific theories, the Church-Turing thesis, the homunculus fallacy, and the misleadingness of the slogan "information is physical." But I learned a lot from you, and you helped me sharpen several points.

Here were some of my favorite pages: 4, 6, 8, 16, 24!, 32, 35-36, 38!, 46!, 47-48!, 52!, 60, 65!, 66, and of course chapter 5. (An exclamation means that something especially intrigued me.)

I know there are a few places where I distanced myself from the phrase "information is physical" in my large Samizdat (`quant-ph/0105039`), but I'm having trouble finding them right now. One is in a letter to Bennett starting at the bottom of page 34. However, the most intriguing moment in my mind is one instance I have not recorded in email before. Let me record it here.

I was visiting Bennett at his weekend home in Wendell, MA one weekend and somewhere in the night I got on a roll about how much I had started disliking the phrase "information is physical." In place of it, I was arguing that a much better, much more accurate call-to-arms we ought to be sending the physics and information-theory communities is that "INFORMATION CARRIERS are physical." Taking that into account is what is behind all the new questions we are asking in quantum information theory. Then I told Bennett, "I am starting to plan a long email that I would like to write to Rolf on this subject. But I know that I'm going to have to word it delicately if I'm going to stand any chance at all of catching his ear and not getting an immediate dismissal as a fool." Charlie replied, "Too late; Rolf just had about a quarter of his brain removed last week." Then he explained that Rolf was just found with cancer in the brain, and that chances were strong he would be dead soon. About a week later Rolf died.

The Shanker 1987 paper looks especially intriguing to me. I'll try to pick it up the next time I'm at Princeton.

By the way, the homunculus fallacy struck me on several levels. I think one might view William James's argument against the correspondence theory of truth in his a little book Pragmatism as a little bit along the same lines, for instance. But it also got me to thinking about one of the things that has long bugged me in the Zurek-style versions of quantum foundations. There, the starting point is how bad the word "measurement" is and how it should be banished from the foundation of the theory. Yet, inevitably (just watch them), whenever push comes to shove in their explanation of the true importance of decoherence, to get the idea across, they start saying things like "in essence the environment 'measures' the system." (Zurek always makes little quote motions with his fingers when he says it.) And that's supposed to lead us to a deeper understanding of that tabooed subject?!?!? (I thought I had put that complaint in a footnote in some recent paper, but for the present I couldn't find it either! I must be losing my memory or losing my mind ... or both!)

Anyway, again, I really enjoyed the thesis. Keep up the good work.

P.S. I also had a look at your webpage. From the schedule of your discussion group, it looks like you're trying hard to make sense of the notions of 'objective chance' and 'propensity' to yourself. I went down that path once—roughly from 1991 to 1996—and it was instructive. I found that I couldn't buy any of the theories, and that's what ultimately pushed me down the long road to Bayesianism. I keep my fingers crossed that you meet the same frustration!!

05 February 2002, to Caves-Schack, "Fighting Windmills"

Thanks for the revealing note on Feller (which I've finally had a chance to read).

Schackcosm 55: *Chris, do you know about any other book or paper that explains a "modern understanding of probability"?*

No I don't really. In fact I don't even really have a strong understanding of where most physicists have gotten their prejudices toward probability. Maybe a good place to start in such a quest is to look at the discussions of probability in the main graduate and undergraduate quantum mechanics textbooks used today (Cohen-Tannoudji or however you spell it, Merzbacher, Liboff, Bohm, and whoever else).

All I remember from my undergrad statistics course was that the professor told us that the Bayesian flavor of statistics was nonsense. Unfortunately, I don't remember the text we used. In fact I got very little out of the course that I remember at all.

06 February 2002, to Caves-Schack, "The Commitments"

Schackcosm 56: *The other problem is that your conversation is far too playful. State assignments are compilations of betting odds. They are COMMITMENTS."*

Dear old friends,

I'm going to try to write a note today that essentially has been sitting in my head for over a month. I'm sorry that I've held on to it for so long, but the great difficulty has been that time has just been stolen from me left and right ever since a few days before Katie's birth. (Katie herself, by the way, is the least of my problems; she's really a dream of a child.) Anyway, I think there were times in this period where the issues were so at the top of my mind that this note could have (or would have) turned out far more passionate and, thus, maybe clearer. But that time has past. Still I'll give it my best shot with the time I can muster today, and also hope that it's not so late as to clash with anything Carl has written for the newest version of the post-BFM paper. Here goes.

In a nutshell, what I'm going to say is that I now take Schackcosm 56 absolutely seriously—indeed probably far more seriously than it was ever meant to be taken.

I used to say that quantum states are "states of knowledge," or "states of information." But, you both know that brooding over the BFM paper caused me to disabuse myself of that. Then I got into the habit of calling quantum states "states of belief" in analogy with what the more left-wing, de-Finetti-flavored Bayesians say of classical probabilities. Now, what I'm going to tell you is that in contemplating the points you *two* think are important in our reply to BFM, I've gotten into the habit of calling quantum states "states of commitment." A quantum state should be viewed most properly as a compendium of commitments, gambling commitments.

Indeed you guys almost nudged me directly there by your love affair (or at least Carl's love affair) with strong coherence. I'll come back to say what I mean by this in much greater detail in a minute, but for the moment let me bring the postulate into a clearing so that you can go ahead and start aiming your arrows at it:

> A quantum state corresponds to a compendium of gambling commitments (i.e., just like the gambling odds of a Dutch-book argument) one is willing to make in various given practical situations. However, the key new point is that these commitments are with respect to ALL THINGS CONSIDERED. There are times when one's commitments correspond to one's (internal) beliefs—and because of that the quantum state remains just as subjective as ever in my mind—but that need not always be the case. There are times (and these are probably the vast majority of all real cases) where the quantum state one ends up ascribing to a system is something less than a compendium of ANYONE's beliefs.

I know that I don't have to remind you that I tried pretty darned hard to choose all my words very carefully in that definition.

Now let me tell you about the sorts of thoughts that lead to this by explicitly replying to some of your old notes to me.

Ruediger wrote:

Schackcosm 57: *I don't quite understand Chris's problem. Strong consistency DOES have a motivation which is very similar to ordinary consistency, and which DOES have the same flavor. It's just stronger. If you violate strong consistency, you are imprudent in the sense that you accept a bet in which, according to your own state belief, you never win, but you lose for at least one outcome that you believe is possible. That qualifies as imprudent, I think. Actually, I believe that the term "imprudent" fits better here than in the case of ordinary consistency, where "outright crazy" seems more appropriate.*

And then Carl wrote:

Cavesism 64: *I agree completely with the above, especially with the distinction between "imprudent" and "outright crazy." I think a violation of strong consistency is somewhere between imprudent and outright crazy, but I haven't been able to think up the right [word]. People who take imprudent actions expect to win big, I think; they are judged imprudent because soberer people can see that the chance of [a] big win is small, whereas the chance of serious losses is large.*

First, let me try to settle the issues of language in these two remarks, or at least try to say more clearly why I was concerned—actually in a rather offhand way in the beginning (see my note of 12/5/01, titled "Dear Prudence")—with the appropriateness of the usual terms, "coherence" and "consistency." Then I'll tackle the more substantial issue Ruediger brings up before that.

I once had an officemate who committed suicide. That is an action I would call (and did call) "outright crazy." Seeing the pain and the soul-searching it put everyone through who was near him, I might even have called it "moronic." And that is important. For, the point I was trying to make explicit to you in my earlier (shoddy) note, is that that kind of craziness is of a much less absolute character than the kind of craziness one would be committing by asserting both A and $\neg A$. In the first case, the craziness is conditioned by one's culture and one's customs, or you might say by the instinct for one's survival. In the second, the craziness is in the breaking of a timeless, Platonic, a priori, ideal "law of thought."

To me, the word "inconsistency"—and therefore the word "consistency"—seems far more to connect to such an ideal Platonic stasis than one's willingness to be taken to the cleaners by a Dutch-bookie. Similarly with the word "incoherent." However, being Dutch-book coherent strikes me much more as an expression of the survival-instinct type than anything else. It is a formal expression of "thou shalt not commit suicide." Why should that be considered an inherently "logical" commandment? Why should it have the right to live in the Platonic realm of Boolean logic? If you ask me, I would say it is probably much more a manifestation of simple Darwinian evolution. We try to stay Dutch-book consistent—it is our ideal of behavior—precisely because of the survival tool it represents for our kind. However, we know by recent experience, that some cultures bask in the idea that, at times, there are reasons to override one's personal aversion to suicide. Do you know of a culture that, at times, finds it useful to override Boolean logic in its mathematical proofs?

When I wrote my original offhand note, I never imagined the backlash I would get from you guys. Nor did I intend to make another proposal to change established nomenclature. I thought I would just simply get a reply of the sort, "Yes, we understand your concerns, but in this case it really is too late to change the nomenclature. A good fraction of the Bayesian community has been using the terms 'coherence' and 'consistency' for almost 70 years. The best you should hope for is an extra paragraph in the paper's appendix assessing why one ought to consider Dutch-book coherence compelling." And I'll still stand by that.

But I actually think there is a point of more substance in this—one that starts to bite much harder when one debates the relative merits of regular Dutch-book coherence versus its stronger cousin (i.e., strong consistency, strict fairness, or what have you). And the new problem is no longer just a problem of language. To say it again, Ruediger writes:

Schackcosm 58 : *Strong consistency DOES have a motivation which is very similar to ordinary consistency, and which DOES have the same flavor. It's just stronger. If you violate strong consistency, you are imprudent in the sense that you accept a bet in which, according to your own state belief, you never win, but you lose for at least one outcome that you believe is possible.*

You don't know how I ALMOST agree with the first sentence of this! In fact, I could agree with it completely if you would just let me change the first instance of "does" to "can."

> Strong consistency CAN have a motivation which is very similar to ordinary consistency, and which DOES have the same flavor.

I said before that strong consistency was of a different flavor than standard consistency, but let me be more careful now—i.e., now that you've helped me sharpen my point. It is not that the flavor is a priori different; it is that in paying the price of strong consistency, you actually get two flavors for the price of one. That is to say, I cannot agree that strong consistency is "just stronger"—it is that and something more.

Here is a very different (and ostensibly much less satisfying) way of motivating strong consistency:

> AXIOM: Whatever an agent's personalistic beliefs, i.e., whatever personalistic probabilities he has managed to write down in his head for some event, when placing a bet (with a Dutch-bookie or otherwise), he MUST place it precisely according to the gambling odds his beliefs afford.

Clearly a special case of this is that when a person believes $p = 1$ for some event, then he MUST bet as if he is certain that it will occur. I.e., he must be strongly consistent.

But you probably ask, "What is wrong with that?"

I will tell you. In doing such a thing, you would throw away loads of freedom, loads of wiggle room that the standard Dutch-book argument generally leaves in your command. In essence you throw away the freedom to concatenate your Dutch-bookie-game commitments with any of a number of other games. Or still another way to say it—though it's not so nice—you throw away your freedom to lie, even when it is in your best, overall, total interest.

An easy case to see this in is the "double agent" Dutch book argument that I wrote you about on 8 Sept. 2001 in a note titled "Negotiation and Compromise" (starting on page 64 of the mini-samizdat). Suppose that Kiki and I draw our money from the same bank account, and that for some given event I ascribe $p = 1$, whereas Kiki ascribes $p = 1/2$. We've talked about our disparity in assignments many times, but I simply cannot convince her to accept the evidence that led me to $p = 1$; she holds her ground, and I know she would do that in any circumstance, come hell or high water. Now suppose I'm later approached by a Dutch-bookie that I am fairly sure has already gambled with Kiki. What should I do? If I don't want to loose my shirt, then I had better not declare my adamant belief in the ultimate occurrence of the event we are betting. Instead, I should adjust my GAMBLING COMMITMENT to agree with Kiki's (sorely) less than optimal one. (Sorely less than optimal from my point of view, that is.)

If on the other hand, you say that in order to play the game of Bayesian probability, I must be strongly Dutch-book consistent, then I no longer have the option to save my own fortunes in this new game. The point is, standard Dutch-book consistency concatenates well with such a further, ancillary game, whereas strong consistency does not.

But this is just a contrived example. Similar, but maybe better, examples can also be found by looking at the old "Keeping the Expert Honest" game (see pages 20–21 in Ph.D. thesis). And I'm sure there are still other more realistic situations that we could come up with if we just tried.

So, what I am saying is Dutch-book consistency is only the tiniest check on one's gambling strategies, and that is a GOOD thing. A very good thing. If one imposes too much structure—apparently almost anything more than simple DB consistency—then one will be left in a lurch in the real world, a world where negotiation and compromise are the keys to survival.

I truly, truly, truly hope you will see the point of this, but I guess I have gained enough experience in the last email war to be prepared for the worst. Here is what I am mainly afraid of. By holding fast to strong consistency as a reasonable addition to standard DB consistency, one ends up in the quantum case with a nice, tightly mathematical looking theorem (almost) in the BFM style. You cannot tell me that at least the Carl among you is not far more attracted to results like that, than the willy-nilly result one is left with if only standard DB consistency is enforced. At the very least, it makes the paper look far more constructive than destructive ... and that's got to be deemed a good thing, right?

What I guess I am saying is that, concerning what Carl wrote:

Cavesism 65: *I, as you know, like strong consistency and hope that in discussions of it, we can separate the a priori reasons for liking or disliking it from reasons based on the conclusions it leads to.*

I believe that I have done that. However, I hope that you will live up to these hopes too! Every bit of the discussion above (excepting the last paragraph) had nothing to do with quantum mechanics. The argument is purely classical and divorced from the BFM issue. In our post-BFM paper, I certainly have no problem whatsoever delineating the whole hierarchy of conditions, regular DB, strong DB, etc. What worries me though is how the present draft gives pretty short shrift to plain old regular DB — the very one that I myself find the most reasonable (for all the reasons above).

Of course, suspecting that you guys will try to wring me out after spouting this blasphemy, I had hoped to come armed with good knowledge of the literature. Unfortunately, as I already expressed way above, time hasn't been on my side these last few weeks. What I was at least able to do though, was run to the Princeton library one day and amass copies of a lot of papers. I ran from one paper to the next, following the citation trail that each gave. Below is the result: My complete collection of Dutch-book papers.

I wish I could say that I had read these, but for me the skim of a title and abstract doesn't count as "reading" a paper ... as it does for XXX, say. Nevertheless, a couple of points did stand out for me. The most important one is that maybe I'm not alone in thinking that strong consistency goes (far) too far. Hacking (1968) below writes, for instance,

> Abner Shimony called it *coherence*; John Kemeny called it *strict fairness*; today many people speak of *strict coherence*. According to Shimony's definition, a set of betting rates on a series of propositions h_i and e_i is strictly incoherent, when "there exists a choice of stakes S_i such that, if X accepts the series of bets at these stakes, then no matter what the actual truth values of h_i and e_i may be, X can at best lose nothing, and in at least one possible eventuality he will suffer a positive loss" De Finetti had a less demanding concept which is called *coherence*. A set of betting rates is incoherent if, no matter what the actual truth values of h_i and e_i may be, X will suffer a positive loss in every eventuality. Logicians usually think that Shimony has improved on de Finetti's concept of coherence, but statisticians, including de Finetti himself, have seldom been persuaded.

After that, he starts writing about nonBoolean algebras, and I haven't had a chance to try to decipher it. But what is really important to me is that this is the first indication I've run across that de Finetti himself thought about strong coherence and then rejected it. I ask, "Why?" I know that Carl has little use for relying on authority—I'm referring to his quick rejection of even wanting to hear de Finetti's opinion—but that is not what is at issue here. de Finetti is someone who thought long and hard about probability from the Bayesian view. I cannot see how it would not be worthwhile to at least hear his arguments. We might save ourselves time, and we might save ourselves from making mistakes that will make us feel foolish.

Then there is a large set of papers debating whether Dutch-book coherence really has to do with "rationality" or rather something else (as I alluded to above). The titles should give those papers away. I've hardly even skimmed those at all, but you can see from above where my present opinion lies. DB consistency defines an expedient for our actions, but it is hardly more rational or logical than that. Skyrms87 in particular, I am told, argues the opposite point of view. If that is a position that is near and dear to your heart, then maybe it would be worth understanding what he has to say.

OK, clearly I'm petering out and it's after midnight now. Soon I won't even be coherent myself; I'm already starting to see signs of it. Let me try to quickly summarize the whole argument in a few sentences, rather than checking and editing all the above to make it more eloquent, more complete, and more connected.

1. DB coherence strikes me as much more a pragmatic requirement than as any rule of rationality (as the law of the excluded middle is).

2. Thus one is more compelled to consider the pragmatic consequences of standard DB versus strong DB.

3. From that, one sees that strong DB is not just more of the same, but carries with it whole new flavors of behavior. In particular it forces us all to be little George Washingtons—"I cannot tell a lie"—when we have a $p = 1$ assignment in mind.

4. Eschewing that, I am forced to divorce our (pragmatic) gambling commitments from our actual beliefs. Our beliefs can be our commitments, but our commitments need not be our beliefs.

5. Thus it is better to say that "probabilities are our gambling commitments, ALL THINGS CONSIDERED." (with apologies to NPR)

6. Quantum states being compendia of probabilities are thus "states of commitment" full stop.

That's the argument. I'm sure this letter is riddled with typos, but I don't want to hold on to it anymore. France Telecom is coming Thursday and I've been tapped to convince them that quantum information is interesting and that our dabbling in it is a little value-added perk they'll get if they stay our customer rather than running away to Alcatel. Can you believe that?

Anyway, maybe Ruediger will be a little happy to see this when he gets into the office tomorrow ... to see that I haven't really abandoned you.

1. Brad Armendt, "Is There a Dutch Book Argument for Probability Kinematics?," Philosophy of Science **47**, 583–588 (1980).

2. Patricia Baillie, "Confirmation and the Dutch Book Argument," British Journal for the Philosophy of Science **24**, 393–397 (1976).

3. David Christensen, "Clever Bookies and Coherent Beliefs," Philosophical Review **C**(2), 229–247 (1991).

4. Barbara Davidson and Robert Pargetter, "In Defence of the Dutch Book Argument," Canadian Journal of Philosophy **15**(3), 405–424 (185).

5. Richard Foley, "Being Knowingly Incoherent," Noûs **26**(2), 181–203 (1992).

6. Bas C. van Fraassen, "Belief and the Will," Journal of Philosophy **81**(5), 235–256 (1984).

7. Ian Hacking, "On Falling Short of Strict Coherence," Philosophy of Science **35**, 284–286 (1968).

8. Frank Jackson and Robert Pargetter, "A Modified Dutch Book Argument," Philosophical Studies **29**, 403–407 (1976).

9. John G. Kemeny, "Fair Bets and Inductive Probabilities," Journal of Symbolic Logic **20**(3), 263–273 (1955).

10. Ralph Kennedy and Charles Chihara, "The Dutch Book Argument: Its Logical Flaws, Its Subjective Sources," Philosophical Studies **36**, 19–33 (1979).

11. R. Sherman Lehman, "On Confirmation and Rational Betting," Journal of Symbolic Logic **20**(3), 251–262 (1955).

12. Abner Shimony, "Coherence and the Axioms of Confirmation," Journal of Symbolic Logic **20**(1), 1–28 (1955).

13. Brian Skyrms, "Coherence," in *Scientific Inquiry in Philosophical Perspective*, edited by Nicholas Rescher (Center for Philosophy of Science, Lanham, MD, 1987), pp. 225–242.

14. Jordan Howard Sobel, "Self-Doubts and Dutch Strategies," Australasian Journal of Philosophy **65**(1), 56–81.

15. Lyle Zynda, "Coherence As an Ideal of Rationality," Synthese **109**, 175–216 (1996).

06 February 2002, to Caves-Schack, "Actually, the Both of You."

This morning I woke up and re-read the long note I sent you last night. In doing that, and looking back at your other notes again, I now think I was too harsh in singling out Carl when I wrote:

> By holding fast to strong consistency as a reasonable addition to standard DB consistency, one ends up in the quantum case with a nice, tightly mathematical looking theorem (almost) in the BFM style. You cannot tell me that at least the Carl among you is not far more attracted to results like that, than the willy-nilly result one is left with if only standard DB consistency is enforced. At the very least, it makes the paper look far more constructive than destructive ... and that's got to be deemed a good thing, right?

That's not fair to him. In particular, both Ruediger and Carl wrote:

Schackcosm 59: *Of course it is nicer to base one's approach on the weaker concept. What we are discussing in the paper, it seems, is reasons for why the weaker concept does not quite give us what we would like. These reasons are secondary, they have nothing to do at all with the Dutch book justification of strong consistency itself.*

Cavesism 66: *Ruediger's two paragraphs nicely illustrate what I was hoping for when I wrote that I "hope that in discussions of it, we can separate the a priori reasons for liking or disliking it from reasons based on the conclusions it leads to." The first paragraph is about the a priori reasons, and the second is about why its conclusions are important.*

Let me focus in particular on Ruediger's sentence, "What we are discussing in the paper, it seems, is reasons for why the weaker concept does not quite give us what we would like." What is it that *we* would like? And, why would *we* like it?

Clearly, I am most happy with the willy-nilly result that standard DB coherence alone gives. I'm glad you guys made that nice and rigorous. For me, it says that B, F, and M were just way off the mark in trying to dictate what various observers MUST ascribe for their quantum states. It tucks nicely with my very first note to them where I registered some protest.

Now, what is nice about the strong coherence version of the argument is that it does at least get us the "if" part of BFM, and thus gives us something solid with which to compare to their campaign. But I suppose, I have always viewed that as a secondary, rather than the primary, point

of our criticism. It's role is simply in that it tells us what we have to ADD to pure (de Finetti flavored) Bayesianism, to get something that comes close to resembling BFM.

So, my apologies to Carl, and my consternation to the both of you!

06 February 2002, to Caves-Schack, "Definition"

Let me go back to the NPR thing briefly while I've got a couple of moments. "All things considered," what does it mean?

I know that your knee-jerk reaction is going to be that it is a hopelessly vague term. However, I'm going to suggest that it is no more and no less vague (and mysterious in its origin), than "belief" was in the first place. It is just broader in scope. In fact, I think it amounts to little more than the sum total beliefs one possesses. (What more could one consider?)

Thus in setting a quantum state, one sets it according not only to what one believes about the system of particular interest ... but also according to the situation one believes he will be encounter in the laboratory, the purposes of the information he will gather and how it will be used, who will share that knowledge so gained, etc., etc. It may even depend upon how the financial markets are doing, which political party happens to be in power, and so forth.

It just takes seriously the idea, that to any quantum system, what we say about it—i.e., what we are openly willing to bet on it—depends upon many things beside the system itself.

Maybe all of this goes back to a conversation I had with Marcus Appleby while I was in Northern Ireland, just after 9/11. I tried to explain my new point of view—that a quantum state is a state of personal belief—and he replied that he though that had the right feeling, but that he had some kind of "unpinpointable" fear that maybe the idea didn't go far enough. He seemed to be saying something to the effect that beliefs can never truly be considered in isolation from other beliefs. I didn't understand his worry at the time, but I think everything I wrote you yesterday starts to pick up on this line. In this regard, by the way, Appleby suggested I read a book by Michael Polanyi, *Personal Knowledge*. (Polanyi was apparently also a chemist of some renown, I believe he said.) Though I haven't had a chance to dig it up yet.

In any case, the COMMITMENTS one (potentially) makes in the sum total of all gambling situations quantify the quantum state. They give it an operational meaning—in terms of its "cash value"—which goes beyond niggling over the details and merits of "beliefs" versus "culture" versus "all relevant things considered," etc.

06 February 2002, to Caves, "The Great Quantum Well"

Cavesism 67: *Not willing to discuss the technical issues much at present because I'm not going to have time to flesh things out for a while. I get a long way by not being interested in Kraus decompositions, which aren't generally convex combinations of unitaries, and I get a good deal farther by insisting on time evolutions, not just single-time stuff. The rest is going to come from having available some operators that rigidify (and thus provide a physical interpretation) for the vectors in Hilbert space, without which the quantum questions (and Hamiltonians) aren't anchored to anything at all.*

I think the last sentence really captures what sets our goals apart. The image I carry around in my mind presently is that the quantum system is very much like an oracle. We ping it, and it provides us with something that we did not have before, something that we could not foresee

... and therein lies its reality. Therein lies its independence from us. Each quantum system is an instantiation of your great quantum well.

But the "interpretation," the "meaning," of the gifts those oracles give us is set from the outside. Completely from the outside. Or at least that's the point of view I'm pursuing.

I just don't see how that line of thought lies anywhere along the lines of your accusation that quantum questions will thus not be "anchored to anything at all."

08 February 2002, to van Fraassen, "Samizdat and Dutch Book"

I was talking to David Mermin a while ago and he mentioned that you had brought my name up in some email and seemed to be aware of some of my papers. (I hope I got the story straight.) Anyway, that piqued my interest: If true, then you are very likely the only philosopher who has ever noticed my existence! I would certainly love to hear your opinion—both con and pro—of my quantum foundational thoughts (forming as they are). Having the critique of a true-blue philosopher would be most useful for steering me to clarity, or even steering me away from the abyss!

My two most close-to-philosophical pieces can be found on the Los Alamos preprint archive: `quant-ph/0106166` and `quant-ph/0105039`. Or, they can both be found at my (almost empty) website `http://netlib.bell-labs.com/who/cafuchs/index.html` with a couple of other pieces of supporting material.

Also I'm in Princeton from time to time making use of your wonderful library system. Maybe I could drop by for a chat? (Sometime in March or later, actually, since I'll be in Japan until essentially then.)

By the way, I was trudging through various Dutch-book arguments recently, and I came across your paper "Belief and the Will," [Journal of Philosophy **81**(5), 235–256 (1984)] in the process. It looks good. I'll let you know if I form an opinion.

11 February 2002, to Mermin, "Oh Modern Wittgenstein"

Concerning points 1 and 2 in your Tractatus Quantico-Philosophicus, read the anecdotes below. [See letter to Chris Timpson, dated 2 February 2002.]

I exclusively use the word "qubit" for the physical system, i.e., the "carrier" of the information, i.e., the object of one's belief, i.e., the oracle the receiver consults at his end of the game. I never use qubit to mean a quantum state (and I don't think I ever have), but I know it's a serious problem in the community.

11 February 2002, to Mermin, "Re: Oh Modern Wittgenstein"

Merminition 46 : *It's the Cbits I'm more concerned with. "bit" – which Charles wants to reserve for the classical physical system clearly has an important abstract meaning as well, relevant to both Qbits and Cbits.*

Aristotle called it matter, the receptacle for accidental properties (presumably some of which could be binary valued). When we turn our attention to a subset of such properties, and are completely ignorant of which of the two is inherent, then we say that we are missing a bit of information.

12 February 2002, to van Fraassen, "The Will to Believe"

Thanks for the note. You warmed my heart with the sentence, "We were all very intrigued with this 'Bayesian' approach to probability in QM." We (i.e., Caves, Schack, and myself, and sometimes Mermin) know that it's all very much at the beginning stages, but things are starting to fall into place so exponentially fast it evermore gains the air of an inevitability. We've now got a load of material that we haven't published yet (and are working very hard to write up), and I in particular have made a strong phase transition in my attitude toward quantum time evolutions (i.e., their level of subjectivity versus objectivity). So things are just flying.

Thanks very much for putting me on the meeting list. But also I hope to meet you before then. In any case, know that I am always, always available on email. (And as one of the documents I advertised in the last note will attest, that is even my preferred means of communication!) So, please feel free to get a dialog going if you wish: I would relish it.

13 February 2002, to Brassard, "Another Kent Paper"

I think "nonlocality" and especially questions about the restrictions enforced on physical theory—in particular, the description of quantum phenomena—due to no-signaling criteria are red herrings. Our brain pulp is better left for deeper matters. You see, in playing the sorts of games that people have been playing along these lines, the very starting point is to take the idea seriously that the quantum state is a physical property, rather than a description of information, knowledge, belief, betting-odds or what have you. And that is an idea I stopped taking seriously a good while ago.

See Section 6 (and its closing paragraph) in my paper `quant-ph/0106166`. By the way, there is a (relatively minor) technical mistake in that section; let's see if you can find it!

13 February 2002, to Caves, "One More for James"

Cavesism 68: *Curt? Emphatically so. Snide? I don't think so. But I can see that my habitual curtness has led to more than its usual negative reaction, so let me go back to square one.*

William James likes to say that all beliefs are "numerically additional" to the reality they take as their target, even "true" beliefs.

Thanks for the note.

I similarly need to automate my Outlook better: maybe that'll be one of the great triumphs of our time together in Brisbane.

16 February 2002, to many, "Some Things Should Not Pass"

Some things should not pass without our best effort to make them indelible. Yesterday, February 15, Kiki and I had to put our golden retriever Wizzy to sleep. He was the most loving and faithful dog either of us had ever had.

As things happen, Wizzy's last day of life marked exactly eight years from the time he first entered my apartment in Albuquerque; it was my second date with Kiki. I had schemed all day about how I might meet her that evening, and the solution was to cook a meal, a large meal. Hopping around the corner from my apartment to hers, I said, "Would you be interested in dinner tonight? I accidentally made too much." She said, "Sure; I was only going to warm up some

potatoes and cheese anyway. I'll be over in about 20 minutes." A couple of minutes later, I got a phone call asking if she could bring her dog with her. I said, "No problem." It was a sweet and touching sight: Wizzy was a dog so insecure at the time, he never left Kiki's side—he didn't sniff around or explore like most any other dog would have; he stayed in bodily contact with her from the moment he entered until the moment he left.

About four months before that night, Kiki had rescued Wizzy from an animal shelter. He must have had a hard life, we surmise, judging from the scar on his head and the fear he had of brooms at the beginning. We'll never really know how old he was, but comparing him to our other golden, Albert, he was probably 12 or 13 years old when he passed.

The day Kiki met him, she had the intention of looking at two dogs before making a decision of which one to take home. Wizzy was the first. When the attendant let him out of the cage, he so leaned his whole body into Kiki and seduced her with his big, loving eyes, that she knew she couldn't put him back.

Wizzy in fact played a predominant role in my meeting Kiki. For some time I had seen her walking him around the neighborhood, and it dawned on me that since no one was ever with her, she had a chance of being single. I waited for my moment, and it came one morning as I walked across a neighborhood park. I introduced myself by going directly for Wizzy. I said, "What a beautiful dog; what's his name?" She told me, and then I asked in a sort of quizzical way, "Is he a purebred?" She said, "Yes." I said, "I don't think so. I have a golden myself, and they don't look very much alike." Why she accepted a date with me a few weeks later remains a mystery.

This morning I broke the news to Emma, and she became sad. She asked, "Where did he go?" I said, "Back to nature." When she's ready to think about it harder, I'll tell her my (presently) favorite metaphor for what happened: Our finite lives are like little drops of water that have parted from the sea. For a small time we have the chance to move around and determine our courses as we please—to leave a trail behind us. But we all eventually run back into the sea. We never stop being; we just become part of something bigger.

Kiki and I put Wizzy into the ground at sunset yesterday, like Egyptians. We gave him his blanket of eight years so that he would never be cold. We gave him his leash so that he could have an infinity of walks, his bone so that he could have an infinity of chews, and his rubber ring so that he could retrieve it for eternity. I told him that I had always known he was a purebred.

17 February 2002, to Bennett, "The Process"

Bennettism 1: *The main worthwhile thing left behind is not your physical remains, nor even their information content (which presumably reemerges as Hawking radiation) but rather the relative state you leave behind in the Church of the Larger Hilbert Space.*

I don't suppose it's ever struck you what an excessive sort of universe that would be. In a way, everything appears twice over. Once in the state and once in the relative state. (If you accept the existence of the universal wave function, the one determines the other uniquely.) What so moved God that he should make two copies of everything? (Redundancy for the purpose of error correction won't do as a reply!)

18 February 2002, to Preskill, "Psychology 101"

Let me reply to some of your points in a way that doesn't reflect their original order.

Preskillism 1 : *In the past I have sensed that you and I differ in how we regard ourselves. I believe that I am just another physical system governed by the same fundamental laws as any other system. You seem to think there is a fundamental distinction between yourself and the system you are observing. To me the Everett view is appealing because it turns away from this egocentrism.*

It's funny, but when I read this, my reaction went in two rather peculiar directions. First I thought, "I wonder if, in the end, the only thing the great quantum foundations struggles will leave behind is a few psychological observations? If so, what a shame." But secondly, I imagined Galileo hoisting me up to the top of the Leaning Tower of Pisa and dropping me off it along with his two famous stones. Even though I cursed and screamed the whole way down, I went "splat" at the same time that they went "thud."

Here's the psychological thought in a little more detail. One of the things that bugs me about the Everett view is what *I* consider *its* extreme egocentrism! Now, how can that be—both of us accusing the other's view as *the* egocentric view? I'll tell you what I think, trying to express the problem from both sides of the fence.

My side gets to go first. What I find egocentric about the Everett point of view is the way it purports to be a means for us little finite beings to get outside the universe and imagine what it is doing as a whole. And what is it doing as a whole? Something fantastic? Something almost undreamable?! Something inexpressible in the words of man?!?! Nope. It's conforming to a scheme some guy dreamed up in the 1950s.

This whole fantastic universe can be boiled down to something representable within one of its most insignificant components—the brain of man. Even toying with that idea, strikes me as an egocentrism beyond belief. The universe makes use of no principle that cannot already be stuffed into the head of an average PhD in physics? The chain of logic that leads to the truth of the four-color theorem (apparently) can't be stuffed into our heads, but the ultimate operating principle for all that "is" and "can be" can?

It's a funny thing: I don't think I've met anyone who would imagine that mathematics will ever come to an end. Or even that it *can* come to an end. There'll always be new axiom sets to play with, new formal structures to write down. But with physics it's a completely different story. People are always wanting to say, "Well we've finally gotten there." Or, "Even though we're not there, we're pretty damned close." It's OK, even condoned, to have Dreams of a Final Theory. From this point of view, all the mathematics yet to come is worthless as far as the essence of the universe goes; the wad was already shot.

You get the point. It's a psychological one, but it's one that I find overwhelmingly powerful. It is that anytime any of us ever has the chutzpah to say, "Here's an ultimate statement about reality," or even a potentially ultimate one, what we're really doing is painting the world in the image of man. We're saying that the measly concepts we've managed to develop up to this point in time fit the world in a way that none of our previous concepts have, that none in the future will ever do better, and, most importantly, we view this not as a statement about ourselves and the situation set by our present evolutionary and intellectual stage, but rather as a property of the universe itself.

Now let me start moving toward the other side of the fence. The question someone like me—someone who has these kinds of blasphemous thoughts—has to ask himself is, how can I ever hope to be a scientist in spite of all this? What can science and all the great achievements it has given rise to in the last 400 years be about if one chooses to suspend one's dreams of a final theory at the very outset? (Or, to tribute Johnny, how can one have law without law?)

I think the solution is in nothing other than holding firmly—absolutely firmly—to the belief that we, the scientific agents, are physical systems in essence and composition no different than

much of the rest of the world. But if we do hold firmly to that—in a way that I do not see the Everettistas holding to it—we have to recognize that what we're doing in the game of science is swimming in the thick middle of things. We're swimming in this undulant sea, and doing our best to keep our heads above the water: All the concepts that arise in a physical theory must be interpreted to do with points of view we can construct from *within* the world.

That is to say, we have to loosen the idea that a physical law is a mirror image of what "is" in the world, and replace it with something that expresses instead how each of us can best cope with and hope to take advantage of the world exterior to ourselves. This, it seems to me, is something that by its very definition can be stuffed into the human brain. The current state of science is our presently best known means for survival. A scientific theory indeed, from this point of view, is yet another expression of Darwinian principles. Scientific theories evolve and survive because the survivors have a kind of staying power that none of the rest of the competition have. Not because they are part of the blueprint of the universe.

The situation of quantum mechanics—I become ever more convinced—illustrates this immersion of the scientific agent in the world more clearly than any physical theory contemplated to date. That is because it tells you you have to strain really hard and strip away most of the theory's operational content, most of its workaday usefulness, to make sense of it as a reflection of "what is" (independent of the agent) and—importantly—you insist on doing that for all the terms in the theory.

I know you're going to find the last sentence debatable, but that is what I see as the danger in the Everett point of view: You are able—or at least purportedly so—to view the universal state as a reflection of something, but at the cost of deleting all the concrete things it was meant to reflect in the first place. What I mean by this is, if we take any concrete situation in quantum mechanics—a system, a measuring device, and some kind of model for the beginning stages of a measurement—we can indeed construct a Church-of-the-Larger-Hilbert space description of it. I'll grant you that. But try to go the other way around without any foreknowledge of the "measurement": Start with the Church, and try to derive from it that a concrete measurement has taken place, and you encounter an embarrassment of riches. You don't know how to identify the valid worlds, etc., etc. (And, if you ask me, invoking decoherence as a cure-all is little more than a statement of faith that some guy from Los Alamos has the all the answers to all the tough questions the rest of us are too lazy to work out.)

So, I myself am left with a view of quantum mechanics for which the main terms in the theory—the quantum states—express nothing more than the gambling commitments I'm willing to make at any moment. When I encounter various other pieces of the world, if I am rational—that is to say, Darwinian-optimal—I should use the stimulations those pieces give me to reevaluate my commitments. This is what quantum state change is about. The REALITY of the world I am dealing with is captured by two things in the present picture: 1) I posit systems with which I find myself having encounters, and 2) I am not able to see in a deterministic fashion the stimulations (call them measurement outcomes, if you like) those systems will give me—something comes into me from the outside that takes me by surprise.

OK, now let me put myself squarely in your pasture. You worry that having those main terms in the theory refer to *my* (or *your*, or Joe Buck's) gambling commitments, is committing a kind of egocentrism. What respectable theory would refer to my particular vices, my desires, my bank account in making its most important statements?

This is going to surprise you now, but I agree with you wholeheartedly. Even enthusiastically so. Where I seem to disagree is that I do not find this a good reason to promote those vices, those commitments to an unearthly realm and call them "states of the universe" (or relative states therein). Instead, it seems to me to be a call to recognize them for what they are and to redouble

our efforts for getting at the real nub of the matter.

Let me try to give you a way of thinking about this that you might respect. What was Einstein's greatest achievement in getting at general relativity? For the purposes of the present exposition, I would say it was in his recognizing that the "gravitational field" one feels in an accelerating elevator is just a coordinate effect—it is something that is induced purely with respect to the description of an observer. In this light, the program of trying to develop general relativity thus boiled down to trying to recognize all the things within gravitational and motional phenomena that should be viewed as consequences of our coordinate choices. Or to use a phrase I've come to like, it was in identifying all the things that can be viewed as "numerically additional" to the observer-free situation which come about purely by bringing the observer (scientific agent, coordinate system, etc.) onto the scene.

Now the point is, that was a really useful process. For in weeding out all the things that can be viewed as "merely" coordinate effects, the fruit left behind could be seen in a clear view for the first time: It was the Riemannian manifold that we call spacetime.

What I dream for in my foundational program for quantum mechanics is something just about like that. Weed out all the terms that have to do with gambling commitments (I used to call it information, knowledge, or belief), and what is left behind will play a role much like Einstein's manifold.

This much of the program, I hope and suspect you will understand even if you are not sympathetic to it. But, I don't know, you might be sympathetic to it. (Especially if I've done a good job above.) However, it is also true that you have rightly suspected some tendencies in me that go further. In particular, in opposition to the picture of general relativity, where reintroducing the coordinate system—i.e., reintroducing the observer—changes nothing about the manifold (it only tells us what kind of sensations the observer will pick up), I do not suspect the same of the quantum world. This is why I recommend to all my friends that they read William James's little article "The Sentiment of Rationality." It sort of sets the right mindset, even though it has nothing to do with quantum mechanics (other than in the efficacy of taking gambles) and goes much further on religion than I myself would go.

Anyway, here I suspect that reintroducing the observer will be more like introducing matter into pure spacetime, rather than simply gridding it off with a coordinate system. "Matter tells spacetime how to curve *when it is there*, and spacetime tells matter how to move *when it is there*." Observers, scientific agents, a necessary part of reality? No. But do they tend to change things once they are on the scene. Yes. Or at least that's the idea.

Does that mean that the scientific agent is something outside of physical law? Well, to give this an answer, you've got to go back and be very careful to use the picture of "physical law" that I built up at the beginning of the essay. What we are "governed" by, God only knows. He's the one, if anyone, who sits outside the physical universe and has a chance to look back at it whenever he pleases. Our task is to build up as good and solid a set of beliefs as we can from within it. In that way, we increase our survival power, and use our spare time to try to bring forth a few progeny of our own. (I used the word "governed," by the way, because you had used it above.)

If Galileo had dropped me from the tower, I feel pretty confident that I would have gone splat.

Aye yi yi, I wrote a lot. That's the dangers of being jetlagged in a foreign country without one's wife and kids. (I'm in Sendai visiting Ozawa.)

I'm going to have to reply to the other points of your note later.

19 February 2002, to Hardy, "Re-Tackle"

I'm in Japan for a couple of weeks at the moment, and I'm finally getting some time put in the 16-hour days again. (Like I used to in the good old days.)

Anyway, I thought I'd tell you, though it is long overdue, I am finally tackling your 5-Axiom paper again. I'm starting to appreciate it much more for sure. If you just weren't so damned non-Bayesian!!! There's a lot of good stuff in it. My main difficulty at the moment is that you have a couple of moves that I know I don't want to allow into my porn: 1) taking mixtures of states (i.e., allowing probabilities of probabilities) as a fundamental step, and 2) invoking extensions to the Church of the LHS.

But I definitely think you are on the right track. And it's probably just a matter of my searching harder for some Bayesian ways of looking at what you've already done. (I'm doing that as we speak.) The most essential things that strike me are 1) the move to column vectors and thinking of measurement as a decomposition of the state, and 2) invoking a relation between K and N. I think those ingredients are definitely here to stay in my mindset. Also I'm warming up to the continuity axiom. I'll try to write you the reasons why soon. (But I've made promises before.)

Anyway, I'm super- looking forward to your stay.

19 February 2002, to Preskill, "Where to Stop?"

You know I've got a million ways of saying why I don't like the Everett interpretation—none of which you find very convincing—but here's a new thought that dawned on me as I was writing my last note to you, and I wonder what you think. Let me try it out on you.

Everett says, "You know Chris, all these silly things you do like leaving measurement as an undefined primitive, etc., will disappear and find a more satisfactory solution if you'll just lay back, relax, and recognize that the quantum states you're working with are really just relative states ... ones derived from the universal wavefunction under one or another decomposition."

I say, "Aha, OK. Then what is this wavefunction of the universe?"

He says, "Well for that, we ought to consult the Hartle-Hawking paper. Here it is: It's $|\psi\rangle$."

Then it dawns on me. How do I know that that state they wrote down isn't just the relative state of our universe with respect to some super-universe? And how do I know that that state is not itself some relative state with respect to some super-super-universe? And so on ad infinitum.

The point is, what principle of science tells you where to stop? None, I'd guess. Is that troubling? I don't know. But it seems a little fishy to me.

Everett tells me, "You've just got to recognize that the wavefunctions you use on a daily basis simply don't have the same ontological status as my universal wavefunction. You might call them 'states of knowledge' in a way. But my universal wavefunction, now that's the real thing; it's here independently of every man, woman, and child." But then I ask, "Well, why does yours get that exalted position? I claim that it itself is a relative state and you can't prove me wrong."

Like I say, I don't know what I think of this yet, but it does strike me as fishy. Once you get into the game of building a Church of the Larger Hilbert Space, who tells you how many pews to put there? That's not something it seems to me you can ever discern from within the universe. It's an article of religion, it seems to me, much like the imagery the appellation seems to provoke.

19 February 2002, to Preskill, "One More Before Lunch"

Here goes. Let's see if I can be brief enough to finish in time. No easy task for me!

Preskillism 2: *Still, I'm flattered by your persistence. Or are you (as I can't help but suspect) slyly recruiting an Everettite who will make a weak case?*

No, I was honest in all that I wrote in the flattering note.

. . . Damn! Didn't make it. It's after lunch now.

What I was going to say was that I was absolutely honest in why I want you to write the sort of paper that you might for the special issue. The point being that if a physicist really, really does find Everett completely adequate to his needs . . . and can argue that it's not a superfluous addition to what he's actually doing when he's doing a calculation, then that would be an interesting datum for the freaky types like me who see it as an ugly picture of the world.

I wasn't slyly recruiting you to make a weak case. But, of course, I actually did have an ulterior motive—something much bigger in my mind than the needs of the special issue—and now that you've forced my hand, I ought to be up front about it. I was banking that if you really did put your heart into making a convincing case for the Everettistas—i.e., the sort of thing that having to write a paper on the subject might draw out of you—then your intellectual honesty would cause you to see how much of the point of view really hasn't been worked out (yet? or maybe ever will be?). I.e., that they have no convincing/relevant argument for the probability rule, that they seem to require a preferred basis, that they seem to require a preferred tensor product structure, that to make sense of two systems, they need to invoke a third, and so on. And when you started to add all those things up, you would also realize that the Everett picture really wasn't much help after all in getting you to the point understanding what measurement is. That's how I was being sly.

19 February 2002, to Preskill, "Re: Psychology"

Preskillism 3: *Sure, scientists are arrogant. That our puny brains can grasp anything about how the world works is a miracle, and I can't pretend to be able to explain it. But I believe it is so.*

It's not the claims of "anything" that worry me so much. It's the claims of "everything." I.e., that our puny brains can grasp everything (in the sense of an "ultimate physics") is the thing that seems implausible to me. And if we can't have that, then we—or more realistically, those who are inclined to do so—ought to be asking what it is we're shooting for.

Sorry I hit a nerve.

20 February 2002, to Wootters, "Out Loud"

Thanks for thinking out loud. I'll just respond to one point.

Woottersism 1: *Let me think out loud for minute here about your note to John. I can think of a* pragmatic *reason for being an Everettista. (At those times when I am particularly attracted to the Everett view, this is what attracts me.) Even if we can't hope to* know *reality, if we can* guess *a model of reality, this guessing helps science progress. What Everett gives us is a guess at ultimate reality. So let's guess that Everett is right, and then work to falsify this guess.*

It's about the guessing part. I had meant to cover that case with the word "potentially" in this paragraph:

> You get the point. It's a psychological one, but it's one that I find overwhelmingly powerful. It is that anytime any of us ever has the chutzpah to say, "Here's an ultimate

> statement about reality," or even a potentially ultimate one, what we're really doing is painting the world in the image of man. We're saying that the measly concepts we've managed to develop up to this point in time fit the world in a way that none of our previous concepts have, that none in the future will ever do better, and, most importantly, we view this not as a statement about ourselves and the situation set by our present evolutionary and intellectual stage, but rather as a property of the universe itself.

That's actually the point you first inspired in me with your aphorism about the dog. There are some things a dog will never understand; there are even questions he can never understand. Why should we expect the evolutionary chain to stop with us? In a way, this cluster of thoughts that I'm starting to think is a rather strong kind of anti-Church-Turing thesis. That is, I think we've gotten into the habit—and Deutsch tried to codify it in his 1985 paper—of thinking that the Church-Turing thesis implies that once you've got a universal machine (people like to say the human brain is one), then you've reached the end of the line. But one should not forget that what Turing was up to in his 1936 paper was to formalize the notion of what is "humanly computable." (This was a point brought home to me by Chris Timpson's excellent undergraduate thesis `http://users.ox.ac.uk/~quee0776/`.)

By the way, not that it matters too much, but I refined one of my paragraphs in the note before I archived it away:

> The situation of quantum mechanics—I become ever more convinced—illustrates this immersion of the scientific agent in the world more clearly than any physical theory contemplated to date. That is because it tells you you have to strain really hard and strip away most of the theory's operational content, most of its workaday usefulness, to make sense of it as a reflection of "what is" (independent of the agent) and—importantly—you insist on doing that for all the terms in the theory.

And that—by the way again—may have been a point also inspired by someone else, namely Schroedinger. Though I haven't been able to completely track its origin in my mind. Somewhere—maybe his 1935 paper in Wheeler and Zurek—he says something like, "understanding quantum mechanics may not require the addition of more variables, but rather taking some of them away."

22 February 2002, to Scudo, "Getting the Mindset"

You asked for some materials to help you get more familiar with the problem I'd like you to work on.

OK, I'm ready to send you some now. And I'll inundate you, but don't let that frighten you. Only try to understand things to the extent that you've got some free time. (And I well expect you may have none at all!)

The first thing to do is read the Brun, Finkelstein, Mermin paper, "How Much State Assignments Can Differ," `quant-ph/0109041`. This is really the thing that started my thoughts off in the present direction, for I completely disagree with them. Their statement is in ultimate conflict with the Bayesian idea of what a quantum state can be. So, understand their argument.

After that, start reading my new samizdat (i.e., underground publication), which I will send you in a separate email. It's quite large (150+ pages), with plenty of repetition, but it mounts an attack on B, F, and M from just about every direction conceivable. Also there is the fact that as time went on, all the issues became clearer and clearer with me, and so I found crisper and crisper

ways of expressing myself. Still reading it (and reading it carefully) might help you get in the right mindset for any number of problems we might be discussing.

The upshot of much of the samizdat is that, for consistency in one's Bayesianism, one must accept that the *assignment* of a POVM (living on a piece of paper) to a measuring device (living in a laboratory) is a subjective judgment at exactly the same level of subjectivity as the quantum-state assignment. Thus one is presented with a cross-roads. Either one accepts pure Bayesianism and gives up the idea that POVMs and quantum time evolutions have objective ascriptions (i.e., gives up the idea that they are independent of the agent assigning them), or one continues with the belief of objective POVMs and time evolutions and adjusts oneself to the idea that probability has to be objective too.

The direction I personally take is that probabilities are subjective, always. They are never objective. Therefore one must make sense of what one means when one speaks of an "unknown POVM" or an "unknown quantum operation." This is where a new kind of de Finetti representation must come in ... and hence your summer work.

But the most important thing for the present is understanding all the motivation leading up to that point. When you are ready to see a sketch of how the theorem ought to go, look at the note of 19 November 2001 to Caves and Schack titled, "A Lot of the Same."

For completeness sake, I will also send you a draft of the paper that Caves, Schack and I are presently constructing to make some of this official. The part that is maybe the most relevant for your education is the appendix on Dutch-book arguments. You might try to understand that. And I'll write you more about that later.

25 February 2002, to Wootters, "A Wonderful Life"

Thanks for the two notes, and wow, thanks for reading the James essay. Your questions were anything but naive. In fact, they were much needed. In trying to answer them, I think I significantly clarified—to myself even!—what I'm hoping to get at. Besides, I certainly don't have a final stand yet; the whole point of view is in the process of formation and questions like yours really help.

I'll do my best to reply to your questions below, and in the process I think I'll finally compose what I've been wanting to say about your "private-world-within-entanglement" musings. At the end of the note, I'll list some of the open questions on my mind. (These are likely to be the naive ones!)

Woottersism 2: *Of course I'm very sympathetic to the perspective you express in this paragraph ... but couldn't one still argue that as a matter of methodology, the tactic of pretending that we can know the whole story has served science well? We make up a model of the world, and this model gives us something to shoot at. We hang on to the model until we have found an explicit flaw in it (other than the flaw of hubris). And then we move on to a new model.*

I find this an interesting question. On the one hand, I think this strategy does work well in advancing science. On the other hand, scientists (and others) are much too prone to accept as true the pragmatic lie that says we can fully understand the world.

Your note to John P. goes some way toward laying out an alternative methodology. You speak of science in Darwinian terms: the most successful theories survive. How then do we proceed as scientists? I suppose the answer is that we still make up theories and test them, but the theories are not tentative descriptions of the world. Rather, theories are schemes for making predictions. But you obviously also want to say that our theories tell us something *about reality, even if they are not descriptions of reality. Moreover, our theories will tell us more about reality if we identify*

and remove from them those aspects that are subjective. So your view of science is not entirely operational. There is realism in the background.

Have I understood you correctly?

Yes there is certainly a kind of realism working in the back of my mind, if what you mean by "realism" is that one can imagine a world which never gives rise to man or sentience of any kind. This, from my view, would be a world without science, for there would be no scientific agents theorizing within it. This is what I mean by realism: That man is not a priori the be-all and end-all of the world. (The qualification "a priori" is important and I'll come back to it later.)

A quick consequence of this view is that I believe I eschew all forms of idealism. Instead, I would say all our evidence for the reality of the world comes from without us, i.e., not from within us. We do not hold evidence for an independent world by holding some kind of transcendental knowledge. Nor do we hold it from the practical and technological successes of our past and present conceptions of the world's essence. It is just the opposite. We believe in a world external to ourselves precisely because we find ourselves getting unpredictable kicks (from the world) all the time. If we could predict everything to the final T as Laplace had wanted us to, it seems to me, we might as well be living a dream.

To maybe put it in an overly poetic and not completely accurate way, the reality of the world is not in what we capture with our theories, but rather in all the stuff we don't. To make this concrete, take quantum mechanics and consider setting up all the equipment necessary to prepare a system in a state Π and to measure some noncommuting observable H. (In a sense, all that equipment is just an extension of ourselves and not so very different in character from a prosthetic hand.) Which eigenstate of H we will end up getting as our outcome, we cannot say. We can draw up some subjective probabilities for the occurrence of the various possibilities, but that's as far as we can go. (Or at least that's what quantum mechanics tells us.) Thus, I would say, in such a quantum measurement we touch the reality of the world in the most essential of ways.

With that said, I now want to be very careful to distance this conception of reality, from what I'm seeking in the foundation game of quantum mechanics. Here's the way I originally put it to John the other day. Let me repeat a good bit of it so that it's at the top of your mind:

> OK, now let me put myself squarely in your pasture. You worry that having those main terms in the theory refer to *my* (or *your*, or Joe Buck's) gambling commitments, is committing a kind of egocentrism. What respectable theory would refer to my particular vices, my desires, my bank account in making its most important statements?
>
> This is going to surprise you now, but I agree with you wholeheartedly. Even enthusiastically so. Where I seem to disagree is that I do not find this a good reason to promote those vices, those commitments to an unearthly realm and call them "states of the universe" (or relative states therein). Instead, it seems to me to be a call to recognize them for what they are and to redouble our efforts for getting at the real nub of the matter. ...
>
> What I dream for in my foundational program for quantum mechanics is something just about like that. Weed out all the terms that have to do with gambling commitments (I used to call it information, knowledge, or belief), and what is left behind will play a role much like Einstein's manifold.
>
> This much of the program, I hope and suspect you will understand even if you are not sympathetic to it. ... However, it is also true that you have rightly suspected some tendencies in me that go further. In particular, in opposition to the picture of general relativity, where reintroducing the coordinate system—i.e., reintroducing the

> observer—changes nothing about the manifold (it only tells us what kind of sensations the observer will pick up), I do not suspect the same of the quantum world. ...
>
> Anyway, here I suspect that reintroducing the observer will be more like introducing matter into pure spacetime, rather than simply gridding it off with a coordinate system. "Matter tells spacetime how to curve *when it is there*, and spacetime tells matter how to move *when it is there*." Observers, scientific agents, a necessary part of reality? No. But do they tend to change things once they are on the scene. Yes. Or at least that's the idea.

From some of my choices of words, I think you probably got the impression that this thing—this structure within quantum mechanics—that I'm hoping to find at the end of the day is meant to be a model of "reality." Or at least our "current best guess" of what reality is. But no, that's not really what I want. And your questions helped make that much clearer to me. Remember, for me, the mark of reality is its indescribability.

What I'm asking for instead is something like what one finds in the old movie, *It's a Wonderful Life*. That is to say, in our scientific theories, we codify some fraction of what we know about manipulating the world and conditionally predicting the phenomena about us. However, suppose we wanted to get at a measure of our place in the world. How would we quantify it, or at least qualify it? That is, how might we ask how important our lives and agential actions are with respect to the theory we ourselves laid out?

Our only tool, of course, is the theory; for it defines the frame for optimal thinking (and imagination) at any given moment. We can only gauge our measure by deleting the free variable that is ourselves and seeing what is left behind. You surely remember what George Bailey found when his guardian angel granted his wish in *It's a Wonderful Life*. He found that his life mattered. So too is what I suspect we will find in quantum mechanics.

But all of that is the sort of thing I won't be able to say in a conference presentation for quite some time. It's the sort of thing that we discussed once before, in the context of some Jamesian quote. It's the underground reason for the philosophy.

At the level of convincing our peers, let me put it to you this way. Within quantum mechanics, there is an invariant piece which is common to all of us by the very fact of our accepting the theory. That is what we are in search of because in some sense—which need not pertain to a realistic conception of a theory's correspondence to nature—it is the core of the theory. It is the single part that we agree upon, even when we agree upon nothing else. In the direction I am seeking to explore, the quantum state is "numerically additional" to that core. (That is, the quantum state is a compendium of Bayesian "beliefs" or "gambling commitments" and is thus susceptible to the type of analysis James gives in his "Sentiment of Rationality." Our particular choice of a quantum state is something extra that we carry into the world.)

I hope that clears up some of the mystery of my thoughts for you—it did for me. Given John's implicit acceptance of the idea that "a true theory is a mirror image of nature," I should not have said in my note that I agreed with him "wholeheartedly." I do not intend for *any* part of the formal structure of quantum mechanics to be a mirror image of nature (in the sense of a proposed final theory). However, I do not intend to give up the reality of our world either.

From my point of view, the only "true" reality that creeps into quantum mechanics is "in the differential"—i.e., in the changes we induce upon our (personal) quantum states for this and that due to any stimuli we give to or take from the outside world. That, however, is a pretty amorphous thing as theoretical entities go. It is little more than what might have been called in older language, the measurement "click."

There is a temptation to go further—to say that the POVM element E_b associated with a

measurement outcome b is itself an element of reality. But I think that has to be resisted at all costs. There are several arguments one can use to show that the *ascription* of a particular POVM to a measurement phenomenon is a subjective judgment at the same level of subjectivity as the quantum state itself. (In fact the two go hand in hand, one cannot support the subjectivity of the quantum state without also taking the subjectivity of the POVM.) Instead, one should view the (theoretical) ascription of a POVM to an actual measurement device as an attempt to set the significance and meaning of the "click" it elicits. Similarly for the Krausian quantum operation associated with the measurement: It describes the subjective judgment we use for updating our quantum-state assignment in the light of the "click." (If you want more details about these arguments, I can forward you some of my old write-ups on the subject.)

So, you probably ask by now, "What does that leave for the core of the theory? Aren't you throwing away absolutely everything?" And the answer is, "No, I don't think so." Let me give you an example of something which I think is left behind. Recall my favorite argument for why the quantum state cannot be an element of reality—it's the Einstein argument I wrote about in Section 3 of my NATO paper. Once I posit a state for a bipartite system, even though by my own admission my actions are purely local, a measurement on one of the systems can toggle the quantum state of the other to a large range of possibilities. Thus, I say that the quantum state of the far-away system cannot be more than my information or the compendium of subjective judgments I'm willing to ascribe to that system.

Notice, however, that in positing the original state, I had to also implicitly posit a tensor-product space for the bipartite system. Let me ask you this: Once this tensor-product space is set, is there any way to toggle one of the factors from afar just as with the quantum state? As far as I can tell there is not. Thus I would say that the Hilbert space of the far-away system is a candidate for part of the theory's core. Well, the Hilbert space—once the choice of a particular quantum state within it is excluded—really carries no substance beyond its dimensionality d. Thus, in a more refined way of speaking, what I really mean to say is that when I posit a quantum system, I am allowed to also posit a characteristic property of it. It is a property that can be captured by a single integer d.

There are some other things which I can argue will be "left behind" in such an analysis, but I don't want to clutter this note too much. Mainly I presented the example above so that I could give you a clearer sense of how I want to draw a distinction between the rawest forms of "reality" (the surprises the world gives us) and the "core of a theory."

It is the core of the theory (along with the theory as a whole) that I am starting to view in Darwinian terms. But don't we have every right to posit that core as a property of the world itself, at least as long as that belief serves us well? This, as you point out, has been the predominant image of what science is about heretofore.

The only answer I can give you is "yes, we can" (just as indeed we have heretofore). So, your point is well-founded. What I am worried about is whether we *should* posit it so. You say that this view has guided science well in the past. But how do you know? In a world with a view that there is no ultimate law, how do you know that we would not be a thousand years more advanced if we had only better appreciated our role as the substratum of our theories? I think it boils down to the difference between an active and passive view of what existence is about. Or maybe the difference between a positive and a negative view.

To make this point, let me try to put things back into the context of regular Darwinian evolution. Consider the word "elephant." Does it denote anything that exists in a kind of timeless sense, in a way that we usually think—or in my case, previously thought—of physical theories as existing? If the concept of an elephant is worthy of treating as a candidate for an element of reality, then so too will a theory's core.

Well, if we have bought into Darwinism in any serious way, then I would say, no, there is nothing particularly timeless about the concept of an elephant. There was once a chance that it might not even arise in the world. The "elephant" is merely a function of the selective pressures that cropped up in our world's particular history. And, ashes to ashes, dust to dust, the poor elephant may eventually disappear from the face of the universe, just like so many species that arose in the course of evolution only to be never discovered by a single archeologist.

But now, contrast the evolution of the elephant with the possible future evolution of the human species. The elephant was an accident pure and simple, from the strictly Darwinian view. But I would be hard pressed to apply pure Darwinism to the future of mankind. The birth of my oldest daughter, for instance was no accident. Her traits were selected based on personal visions that both her mother and I had for the future. Similarly, but not so excitingly, with the golden retriever, and all our other domesticated species. The key point is that in the present stage of evolutionary development, we have it within our power to move beyond strict Darwinism. This is what our industry of genetic engineering is all about.

However, we would have never gotten to this stage if we had not first realized that the concept of a species is not immutable. As strange—and as crazy and as scary—as it may sound, this is where my thoughts are starting to roam with physical theories. This does not mean, however, that we can have exactly what we want with our physical theories—that they themselves are little more than dreams. Just as the genetic engineer can make a million viruses that will never have a chance of surviving on their own, there is more to the story than our whims and fancies: There is the ever-present selective pressure from the outside. But that does not delete the genetic engineer's ability to make something that was never here before.

But now, I go far, far, far beyond what I needed to say to answer all your questions. Mainly, I just wanted to emphasize why I intentionally placed the words "a priori" in my definition of reality way above.

I fear now slightly that you're going to realize I'm one of the craziest people you've ever met! And, trust me, I'm not sure I really believe all that I said in the last three paragraphs. But it does strike me as a productive, or at least hopeful, train of thought that someone ought to explore. I guess I offer myself as the sacrifice.

— — — — —

There. I think that's enough of my going around your questions in a rather wide way. Let me now zoom back to the center of one of them for purposes of a final emphasis.

Woottersism 3: *But you obviously also want to say that our theories tell us* something *about reality, even if they are not descriptions of reality.*

I hope you can glean from all the above that I do indeed believe our theories tell us something about reality. But that something is much like what the elephant tells us about reality. It's presence tells us something about the accumulated selective pressures that have arisen up to the present date. A theory to some extent is a statement of history. It is also a statement of our limitations with respect to all the pressures yet seen, or—more carefully—a statement of our limitations with respect to our imaginations for classifying all that we've yet seen. (I for instance, cannot jump off the leaning tower of Pisa unprotected and hope to live; you, for instance, cannot get into your car and hope to push on the accelerator until you are traveling beyond the speed of light.) Finally, to the extent that we the theory users are part of nature, the theory also tells us something about nature in that way.

But for any theory, there is always something outside of it. Or at least that's the idea I'm trying to build.

PS. Way above, I said I would finally say a few words about your "private-world-within-entanglement" musings. But somehow it didn't quite fit in with the flow of the rest of what I wanted to say. So, let me try to present the statement in isolation. From my point of view, the quantum state, and with it entanglement, never pierces into the quantum system for which we posit a parameter d (the "dimension"). Similarly for any bipartite system for which we posit two parameters d_1 and d_2. The quantum state is only about what I'm willing to bet will be the consequences when I reach out and touch a system. Otherwise, indeed, a quantum system denotes a private world unto itself. And similarly with bipartite systems. We have very little right to say much of anything about the goings-on of their insides. (This part of the picture is something I've held firmly for a long time; it even shows up in my Physics Today article with Asher.) Thus, I guess what I'm saying is that I can find a way of agreeing with what you wrote me:

Woottersism 4: *Or maybe there* is *a possibility of testing the above distinction, at least in principle. (It depends on how I finish my interpretation.) Let's consider David Deutsch's thought experiment in which a sentient computer measures, in a reversible way, some property that doesn't have a definite value for the object being measured. Later, the computer will report that he observed a definite outcome but did not record which one, and the various possible paths will have been brought together successfully to exhibit interference. That's David's prediction. I would agree that interference should be possible, but I may disagree about what the computer will report. If it is true that one needs to have a* shared *world in order to have a sensible world, then when the computer was all alone in his private entanglement with the object he measured, who knows what he experienced? Maybe he experienced nothing more about that measured object than what one electron in a singlet pair experiences about the other electron. Maybe he cannot honestly report later that he saw one outcome or another.*

But maybe I'm coming at it from a different point of view.

PPS. I also promised to end with some open questions. But I'm petered out now. And if you've gotten this far, you're probably exhausted too. So I'll just leave it for the future, depending upon how interesting you find the ideas above, or how much you think they're nonsense!

Only two and half days left in Japan.

27 February 2002, to Mermin, "An Even Tireder Old Man"

I forgot to include you as a recipient of the note below as I had promised. Let me not tell a lie: It's attached this time around.

I made good progress on several fronts in Japan. Maybe the most important—though most trivial—was that I finally found a rather natural Bayesian motivation for the linearity of the probability rule in Lucien's system. (Lucien is a trickster; he sneaks axioms in left and right that he doesn't call axioms.) Anyway, at the very beginning of the paper already, I couldn't accept his motivation for linearity, falling back as it does on the idea that a state refers in an objective way to (a class of) preparations. That is, he didn't take states to be Bayesian beliefs pure and true, and that's crucial to me.

Here's the solution; it's as simple as can be (now that I see it). The upshot of Lucien's formalism is that a measurement is any "I-know-not-what" that causes a refinement of one's initial beliefs (consistent with the other axioms). Thus a measurement is simply an application of Bayes rule by its very definition—moreover, any application whatsoever within the allowed range. OK, that's good enough for the single observer and already thrills me (because it expresses in a more rigorous way a claim I made in the NATO paper at the level of density operators).

But when there are two observers and you view measurement as nothing but an application of Bayes rule, you have to have some method of saying when they agree that it is the same measurement (so that you don't fall into the infinite regress I tried to warn you of when you were going full-steam-ahead with BFM). Well, here again the solution is simple: Two observers should call a refinement of their beliefs the same measurement, precisely when they draw the same meaning for the data they obtain. Thus as long as both observers ascribe the same $P(\text{data}|\text{hypothesis})$, we will say they "agree" upon the measurement.

And that's it. That does the trick. That and that alone gives linearity. No Gleason's theorem; not even anything fancy. I'll make this prediction right now: Quantum mechanics will turn out to be one of the simplest structures we've ever seen in physics, and for 75 years we've just been too pigheaded to see it.

Somewhere over the pacific,

02 March 2002, to Caves, "Re: Mermin Festschrift"

Cavesism 69: *I'm inclined to think that the Mermin festschrift would be just the right place for my speculations about the reality of Hamiltonians, especially since he has kind of encouraged this idea. What do you think?*

I think, "I wonder why you ask *me* this!?!?" Are you contemplating that we might have dual submissions? Abbott and Costello, Crosby and Hope, Lewis and Martin, the Smothers Brothers? The problem is, both of us would want to be the straight man.

OK, enough sarcasm. To answer your question, yes, that probably would be a good idea.

By the way, here's a little technical meat for the hungry wolf. Suppose I were willing to go with you on the objectivity of THE interaction Hamiltonian or unitary in a measurement interaction. (I'm not willing in actual fact, but that's beside the point for this question.) If we write down the POVM such an interaction ultimately leads to, then we will get something like this for its elements E_b:

$$E_b = \mathrm{tr}_{\mathrm{A}}\Big((I \otimes \sigma)U(I \otimes \Pi_b)U^\dagger\Big) \ ,$$

where all the terms have the usual interpretation. I believe, according to you, there should be only one subjective piece to the right hand side of this equation, namely the ancilla's initial density operator σ.

Here's my (innocent, first) question. The subjectivity in σ will afford a range in our interpretation of which POVM was actually measured in any such interaction. What is that range? How large is it?

A corollary question—this one is more rhetorical and less innocent—is this. You know I have never liked the above formulation of what a measurement is because one has to invoke a second measurement (the Π_b) to explain the first (the E_b). This leads to an infinite regress because one can ask (as von Neumann did), where do the Π_b come from. Thus, just to make sure we're on board, let me reconfirm that I got the setting in your mind right in my elaboration above. Namely, that the only truly subjective term you would accept in the displayed equation is σ? The range in E_b is consequent only to that? If that's the case, I would also like to understand how you invoke a stopper at the level of the Π_b. That is, why do the Π_b not have some of their own range, induced by a higher level measurement model?

You can tell I'm still deeply jetlagged.

03 March 2002, to Scudo, "De Finetti and Strong Coherence"

Thanks for the ghost story.

Scudoism 1: *Will tell you more as soon as I finish my calculations for this paper.*

Not to worry at all. I have a hard enough time leading a rushed life. You should not expect that I expect that out of you when I cannot live up to standard myself! When you arrive here you'll have plenty of time to practice the macabre.

But still, if you've got the interest, who am I to get in your way. Here's some historical investigative work that you might tackle if you find that have nothing better to do. In fact, your friend Regazzini might be a wonderful resource in this regard. In the draft of the anti-BFM paper that Caves, Schack, and I are putting together, we make a distinction between simple "Dutch-book consistency or coherence" (i.e., the notion that de Finetti and Ramsey first introduced) and "strong Dutch-book consistency." The latter is a notion that apparently Abner Shimony first introduced, though Caves and Schack stumbled across it independently in one of our many wars. (See reference 12 on page 139 of the samizdat I sent you for the original citation.)

Starting on page 133 of the samizdat, I write a fairly strong polemic against the "requirement" of strong-consistency—it seems to me that regular Dutch-book consistency is the most we can demand. Here's my question. As you'll see from the quote of Ian Hacking on page 137, de Finetti was aware of Shimony's addition to coherence and apparently rejected it. Do you think it might be possible to find out where he wrote about this? More importantly, what were his particular reasons for rejecting the stronger notion?

04 March 2002, to Folse, "Sliding Off the Deep"

I was able to print out your paper without a hitch. And thanks moreover, for your detailed notes on my letter to John Preskill. Let me try to answer some of your questions (in particular the one about Popper and propensities) by sending you still more rubbish. It's in the form of a follow-on letter (to the Preskill one) that I sent Bill Wootters while still in Japan. Of course, I would appreciate any further comments like your last ones if you've got the time!

Let me quickly reply to two of your points explicitly.

Folse-ism 1: *Do any physicists still think this way? I realize cosmologists talk about a TOE, but surely the conduct of contemporary research is not animated by the thought that we're closing in on some "final" description of the universe? That was true enough a century ago, but I think that the weight of philosophy of science at least in the 20th century has pulled against that sort of image of science – certainly in a post-Kuhnian era. It is curious perhaps that this kind of attitude might still persist in so-called "foundational" studies, whereas I would suppose in something like biological research everyone would have become much more historicized by now.*

Nope, the attitude runs pretty rampant in the quantum information and computing community. (I'm definitely one of the outsiders there.) Have a look at David Deutsch's *The Fabric of Reality*. I think he expresses much of the majority opinion of our little clique there. But I think you'd be surprised to know that the "dreams of a final theory" attitude runs pretty rampant even in such workaday fields as laser physics. I just got the following message from an experimentalist friend of mine whose work is predominantly used for the design of better fiber optics:

El Jeffy 4: *One of the thoughts that continues to strike me is your optimism in the continual evolution in physical theory. I will admit that on bad days I feel like we are in the twilight of physics - only incremental progress is left to be had.*

Of the physicists I have met who have even heard of Kuhn or Rorty—there aren't that many—almost all of them view what little they know of their thoughts (i.e., Kuhn and Rorty's) with a little contempt. (By the way, John Preskill's reaction to my letter was particularly violent. So I think it is the ideas, and not the men, that bug their sensibilities.)

By the way, I have indeed just discovered Rorty and I find him wonderful. I've read the two collections of essays *Philosophy and Social Hope* and *Consequences of Pragmatism*. It won't be long, and I'll move on to the book of his that you recommended. Also since Vaxjo, I've gotten a good feel for James (I read *Pragmatism*, *The Meaning of Truth*, and Perry's massive biography of him), and also of Dewey (though I've only read *John Dewey: The Essential Writings*).

Thanks for telling me about Kitcher; I hadn't heard of him. And thanks for the pointer to Toulmin; I'll look more deeply into him than I have in the past.

Concerning one of your other points,

Folse-ism 2: *It is arrogant perhaps, but I don't see this as the egocentricity. Every attempt to sketch a conception of the universe from our best theories at any date in human history in effect commits such arrogance. Were the Newtonians of the end of the seventeenth century being "egocentric" to think that Sir Isaac had done nothing less than peer into the mind of the Divine and discerned God's blueprints for the universe?*

Yes. (In my opinion.) And you might interpret James and pragmatism in general as a reaction to that. However, I think in our modern age with quantum mechanics we have a motivation and opportunity in front of us that James did not have. Try to give quantum mechanics a naive realist interpretation—you can do it, or at least both Everett and Bohm tell us we can—and you find yourself contorting yourself beyond belief. It's as if nature is telling us for the first time, "Please don't interpret me in a naive realist fashion. I can't stop you, but please don't."

OK, the sun is rising and the vampire must return to his native soil.

04 March 2002, to Nicholson, "The Good Questions of Nicholson"

El Jeffy 5: *So. My question to you is two-fold. Do you still feel the same way about particle physics experiments that try to continually increase the energy of interactions between particles? And, if not, what extrema in physical theory should we be poking at in order to continue the evolutionary progress?*

You ask a good question old friend—the second one—and I shall have to think a while before I reply. Or at least compose myself before I reply.

In the case of your first question the answer is easy. Yes, I still feel that way. And the feeling is not that the SSC would not have been worthwhile science per se; it might have been. It was just the cost versus insight-and-control-of-nature ratio that doomed it in my mind. Furthermore, the paradigm the high-energy types presuppose certainly puts me off. They think they are doing something deeper than tabulating the Hamiltonian of an iodine atom, say, but I don't think they are. They are just tabulating the Hamiltonian of another system, and that has the value it always has.

But let me think harder about how to answer your second question. For my present project (i.e., quantum mechanics), it strikes me not so much as poking at the extrema that we ought to be spending our time on, but in poking at the consistency of the worldview we wish to embed the theory in. To revert to my analogy of the construction of general relativity, it was only consistency that Einstein was seeking when he first started contemplating how gravitation fit within the framework of special relativity. Seeing that it did not (and weeding out all the coordinate effects as I expressed earlier) is what led to the bending of the manifold.

But don't take that as my definitive (or exhaustive) reply. I'm still thinking.

06 March 2002, to Hardy, "Poetry on Concrete"

I decided to write this letter in LaTeX so that I would have a small chance of being clearer. A few months ago, you wrote me this:

Hardyism 1: *[M]y basic picture of preparation/transformation/measurement is rather similar, though considerably less poetic, than your picture of humans interacting with a world. ... After the poetry you need to make the journey from vision to concrete construction.*

You're quite right. I think your paper has finally started to mix the right sort of concrete to support a building like the one I want to see built.

In what follows, all I'm going to do is keep true to this word I learned—poetaster—and try to write a little amateur poetry at your construction site. There are three things that really intrigue me in your paper, and one that really confuses me. I'll take them one at a time.

The first is simply the absolutely beautiful ease with which you bring us to a vector space structure. I had known about the concept of "informationally complete" POVMs for years[2]—they're even a crucial part of the proof of our quantum de Finetti theorem[3]—but before your work I had never appreciated so clearly that the probability distributions they generate ought to be taken as the very most primordial representation of a quantum state.

> *What is a quantum state? Nothing more than one's gambling commitments with respect to the outcomes of a fiducial measurement. Full stop.*

Or at least, that's the way I want to view it.[4]

Now, what you write presently is not completely consistent with that characterization. I'm going to do my best to try to make that plain to you. You write, "I don't believe that my approach needs to adopt a particular approach to interpreting probability." There is a sense in which that is certainly true for the mathematics you've already performed. However, if we're ever going to get past this foundational impasse in quantum mechanics, I would say firmly that we cannot leave the story there.

Here is the difficulty. If you take an objective approach to probability, then you are necessarily left with the conclusion that a measurement "here" in general changes something objective (or physical) over "there." And you will never get around that. Just consider any entangled state. Making a measurement on Alice's side changes her ascription—without any considerations of time

[2] I think the idea can be traced back to E. Prugovečki, "Information-Theoretical Aspects of Quantum Measurements," Int. J. Theo. Phys. **16**, 321 (1977).

[3] See Section IV of `quant-ph/0104088`.

[4] And I say this even though I was saying things right after the Växjö meeting like I said to Khrennikov in the letter of 4 July 2001 attached below. Somehow, I just never took the idea so seriously as I do now, i.e., after reapproaching your paper in Japan. "Fiducial measurement" + "a notion of transformation from one measurement to another that carries the probabilities with it" is what tells the whole story!

or distance—of the state on Bob's side. If a state is purely equivalent to a set of probabilities and probabilities are not subjective degrees of belief or gambling commitments, but rather objective and independent of the agent, then something physical must change at Bob's site with that far away measurement. And if you believe that, then you end up with conundrums out the wazoo.

I know that may not be your stand yet, but it is mine. Thus I'm left with the task of putting your mathematics into a framework I can accept. The first and clearest example of a technical mismatch between our views is what we each might consider as a valid motivation for the convex-linearity rule in your Eqs. (41)–(43).[5]

I think if we were to trace the roots of our mismatch, we would find it in that I do not subscribe to your "basic picture of preparation/transformation/measurement" ... which I think you think is common to both of us. Here's the way I put it to Asher Peres when we were constructing our Physics Today paper:

> In general I have noticed in this manuscript that you lean more heavily on the word "preparation" than we did in our letter to Benka. ... Unless I misunderstand your usage of the word, it may actually be a little too anthropocentric even for my tastes. The problem is this: consider what you wrote in the paragraph about the wave function of the universe. It seems hard to me to imagine the wave function of those degrees of freedom which we describe quantum mechanically as corresponding to a "preparation." Who was the preparer?
>
> It is for this reason that Carl Caves and I prefer to associate a quantum state (either pure or mixed) solely with the compendium of probabilities it generates, via the Born rule, for the outcomes of all potential measurements. And then we leave it at that. Knowing the preparation of a system (or the equivalence class to which it belongs) is one way of getting at a set of such probabilities. But there are other ways which surely have almost nothing to do with a preparation. An example comes about in quantum statistical mechanics: when the expected energy of a system is the only thing known, the principle of maximum entropy is invoked in order to assign a density operator to the system. There may be someone beside me in the background who knows the precise preparation of the system, but that does not matter as far as I am concerned—my compendium of probabilities for the outcomes of all measurements are still calculated from the MaxEnt density operator.
>
> To help ensure that I was not jumping to conclusions on your usage of the term, I reread today your paper "What Is a State Vector?" [AJP **52** (1984) 644–650]. There was a time when I agreed with everything you wrote there .... But as of today at least, I think a more neutral language as in our letter to Benka is more appropriate.

Thus what I seek is a picture that involves only measurement and transformation. (And if the truth be known, I'd like to get rid of transformation *in a sense*, leaving only measurement behind from your trinity. But that's a longer story, and I'll only give hints of it below.) However, I do think your formalism is just about up to the task, despite your Figure 1 on page 4.

The way I see it, this thing called the "preparation" is just a token for the *right and true* quantum state you imagine for a system. But from a steadfast-Bayesian[6] point of view, there is no such thing as a right and true probability assignment for anything. Consequently, if a quantum

[5]By the way, I have to take a little exception with your accounting rules for your axioms. In similar treatments from the "operational school," say of Kraus, Holevo, and Ozawa, a convex-linear or affine assumption (of exactly the same spirit as your Eqs. (41)–(43)) is always listed as a distinct axiom.

[6]Some would say "radical Bayesian."

state is to be solely a probability assignment for the outcomes of a fiducial measurement, from the steadfast-Bayesian view, there can also be no such thing as a right and true quantum state.[7]

Let me emphasize this once more at the purely classical level before going on. For the Bayesian there is no such thing as a *right* probability assignment; there is no such thing as a *wrong* probability assignment. A probability assignment is a set of numbers that signifies which gambles one is willing to make, period. How those numbers are set in any given instance is an issue that has nothing to do with nature. Instead it has everything to do with all the ugly things we try to keep out of science: one's emotions, one's intelligence, one's hopes, the rumors one has heard recently, and so on. In fact, a probability assignment has nothing to do with the world itself—in a sense it floats above the physical world.

What this boils down to is that—*within a Bayesian paradigm*—your argument for convex-linearity cannot make any sense:

> Assume that the preparation device is in the hands of Alice. She can decide randomly to prepare a state $\mathbf{p}_A$ with probability λ or a state $\mathbf{p}_B$ with probability $1-\lambda$. Assume that she records this choice but does not tell the person, Bob say, performing the measurement. Let the state corresponding to this preparation be $\mathbf{p_C}$. Then the probability Bob measures will be the convex combination of the two cases, namely
>
> $$f(\mathbf{p}_C) = \lambda f(\mathbf{p}_A) + (1-\lambda) f(\mathbf{p}_B) \tag{1}$$
>
> This is clear since Alice could subsequently reveal which state she had prepared for each event in the ensemble providing two sub-ensembles. Bob could then check his data was consistent for each subensemble. By Axiom 1, the probability measured for each subensemble must be the same as that which would have been measured for any similarly prepared ensemble and hence (41) follows.

Here's the difficulty. What does the distribution $(\lambda, 1-\lambda)$ refer to? From the Bayesian view, it can only refer to a belief (or better yet, a gambling commitment) that Bob possesses concerning Alice's actions. On the other hand, the numbers $f(\mathbf{p}_A)$ and $f(\mathbf{p}_B)$ refer to a couple of alternate beliefs a completely different person, Alice, might possess. However—*from the Bayesian standpoint*—there is nothing *a priori* that can be used to rigidly identify Bob's beliefs as a function of Alice's ... as Eq. (1) above would imply. Alice has every right to believe anything she wants; Bob has every right to believe anything he wants.[8]

For the Bayesian, all the action—all the science—is not in the probability assignments themselves that various individuals might make, but in how those assignments change upon the acquisition of new data (steered by the influence of rationality).[9] In particular, the important things to ferret out are the conditions in any situation where two observers will converge in their beliefs

[7] I have argued this point at great length in a correspondence with Caves, Schack, and Mermin predominantly. I'm going to post that on my website `http://netlib.bell-labs.com/who/cafuchs/index.html` at just about the same time I send this note off to you. The file is called "Quantum States: What the Hell Are They?," and the relevant part with respect to my remark starts on page 19 and winds up somewhere around page 97. However, the most relevant notes are the ones to Brun-Finkelstein-Mermin dated 7 August 2001, to Caves-Schack dated 22 August 2001, and to Mermin dated 2 September 2001. I hope you will have a look at those at the same time as reading the present note.

[8] Nor does it help to try to divert the discussion from a two-player situation and stuff all of these beliefs into a single head (Alice's say). For at any instance, Alice believes what she believes. And there is nothing within the Bayesian creed to force what she ought to believe in various other hypothetical situations—for example, where the distribution $(\lambda, 1-\lambda)$ might describe the amount of amnesia she has concerning a previous belief that she's just forgotten.

[9] Read my note to Mermin titled "Reality in the Differential" starting on page 128 of my previously mentioned posting.

when they agree that they are acquiring the same data. Bernardo and Smith put it in a beautiful way:

> The subjectivist, operationalist viewpoint has led us to the conclusion that, if we aspire to quantitative coherence, individual degrees of belief, expressed as probabilities, are inescapably the starting point for descriptions of uncertainty. There can be no theories without theoreticians; no learning without learners; in general, no science without scientists. It follows that learning processes, whatever their particular concerns and fashions at any given point in time, are necessarily reasoning processes which take place in the minds of individuals. To be sure, the object of attention and interest may well be an assumed external, objective reality: but the actuality of the learning process consists in the evolution of individual, subjective beliefs about that reality. However, it is important to emphasize, as in our earlier discussion in Section 2.8, that the primitive and fundamental notions of *individual* preference and belief will typically provide the starting point for *interpersonal* communication and reporting processes. In what follows, both here, and more particularly in Chapter 5, we shall therefore often be concerned to identify and examine features of the individual learning process which relate to interpersonal issues, such as the conditions under which an approximate consensus of beliefs might occur in a population of individuals.
>
> — pp. 165–166, Bernardo and Smith, *Bayesian Theory*

> What is the nature and scope of Bayesian Statistics within this spectrum of activity? Bayesian Statistics offers a rationalist theory of personalistic beliefs in contexts of uncertainty, with the central aim of characterising how an individual should act in order to avoid certain kinds of undesirable behavioural inconsistencies. The theory establishes that expected utility maximization provides the basis for rational decision making and that Bayes' theorem provides the key to the ways in which beliefs should fit together in the light of changing evidence. The goal, in effect, is to establish rules and procedures for individuals concerned with disciplined uncertainty accounting. The theory is not descriptive, in the sense of claiming to model actual behaviour. Rather, it is prescriptive, in the sense of saying "if you wish to avoid the possibility of these undesirable consequences you must act in the following way."
>
> — p. 4, Bernardo and Smith, *Bayesian Theory*

And, it is precisely this that I'm going to use to get to your Eq. (41) from a Bayesian starting point. But before I can do that, let me praise the second thing that I see as a deep contribution of your paper. This is the very *definition* of what a measurement is: It is *any* application of Bayes' rule consistent with your remaining axioms.[10]

To make some sense of this, let me put the problem into a notation that is slightly better for my purposes. Suppose the outcomes of a fiducial measurement are labeled by h and the outcomes of some other measurement we might contemplate are labeled by d. Then I will dually write the quantum state (via your identification of states with probabilities for the outcomes of a fiducial measurement) as a function $P(h)$ or as a vector $\mathbf{P}$. What I mean by an application of Bayes' rule

[10] There is no doubt that I have been predisposed to saying something like this for a long time. For instance, at the top of page 26 of my NATO paper `quant-ph/0106166`, I say, "Quantum measurement is nothing more, and nothing less, than a refinement and a readjustment of one's state of knowledge." The difference is, your paper brings this beyond just some metaphor. There I tried to capture the idea with my Eqs. (57)–(59), which are explicitly in terms of density operators. You on the other hand, do it with the real thing. I.e., you make it absolutely airtight that it is Bayes' rule operating in the background, and not just some noncommutative analog of it.

is the supposition of a set of probability distributions $P(h|d)$ (or $\mathbf{P}_d$ in vector notation)—one for each d—such that

$$P(h) = \sum_d f_{\mathbf{P}}(d)\, P(h|d)\ , \tag{2}$$

or alternatively

$$\mathbf{P} = \sum_d f_{\mathbf{P}}(d)\, \mathbf{P}_d\ , \tag{3}$$

for some other probability distribution of $f_{\mathbf{P}}(d)$ over the variable d. When a particular data d is collected, one—at least tentatively—enacts the transition

$$\mathbf{P} \ \longrightarrow\ \mathbf{P}_d\ . \tag{4}$$

This is what I claim you have taught us is "what a quantum measurement is."[11]

A measurement is any "*I know not what*" that enacts a transition of the form Eq. (4). What is the variable d? How is it defined? It simply is not, except as a place holder in a particular instance of Bayes' rule. And that's it: That's the story of measurement.

Now, where does one get the convex-linearity of the probability rule $f_{\mathbf{P}}(d)$ in a Bayesian-happy way? The technique is to consider when it is that two observers will think they are performing the same measurement. That is to say, I could walk into the room and think to myself that the measurement device in front of me gives me warrant for the decomposition in Eq. (2). You on the other hand, with a completely different set of beliefs and experiences, may think that the device warrants you to the decomposition

$$Q(h) = \sum_d f_{\mathbf{Q}}(d)\, Q(h|d)\ . \tag{5}$$

When shall we say that we actually agree upon the measurement? Classically, the answer is when the statistical model made use of by each of us is the same:

$$Q(d|h) \equiv \frac{f_{\mathbf{Q}}(d)\, Q(h|d)}{Q(h)} = \frac{f_{\mathbf{P}}(d)\, P(h|d)}{P(h)} \equiv P(d|h)\ . \tag{6}$$

That is, if we had hold of which one of the fiducial outcomes actually obtained, then we would draw the same meaning from it. Here, by "meaning" I mean how much we would feel warranted in revising our beliefs that a d would have popped up if we had instead performed an appropriate measurement for it.

By the way, notice one thing. If this account deviates from standard Bayesianism any at all, it is only in this. Never once did I invoke the *necessity* of a joint probability distribution

$$P(h,d) \equiv f_{\mathbf{P}}(d)\, P(h|d) \tag{7}$$

in my description of Bayes' rule. Of course, such a function exists as a mathematical artifice—i.e., it has all the properties of a joint probability distribution—but one should not try to make any meaning for it beyond that. In particular—and, especially with regards to the quantum context—one should not feel it necessary to interpret the function as a degree of belief of the simultaneous *existence* of a true h and a true d.

[11] Note carefully that I used the word "tentatively" in this description. That is because there is a possibility of the further quantum phenomenon that measurement can be more invasive than supposed classically. What this means operationally is that the $\mathbf{P}_d$ may not ultimately concern the original fiducial measurement, but may instead be about a completely different fiducial measurement at the end of the process. See my discussion centered around Eqs. (57)–(59) of `quant-ph/0106166`.

OK, let me use that now to start talking about convex-linearity again. Suppose there are at least three agents on the scene, Alice, Bob, and Charlie. And suppose Alice subscribes to Eq. (2) for her description of what the device in front of her is about, whereas Bob subscribes to Eq. (5), and Charlie still further subscribes to,

$$R(h) = \sum_d f_{\mathbf{R}}(d)\, R(h|d) \; . \tag{8}$$

If that is as far as it goes, then there is nothing whatsoever we can say about the relation between $f_{\mathbf{P}}(d)$, $f_{\mathbf{Q}}(d)$, and $f_{\mathbf{R}}(d)$. Indeed there need be no relation whatsoever.

However, suppose Alice, Bob, and Charlie share the minimal belief that actually the same measurement is being performed with respect to each of their descriptions. Then by definition

$$P(d|h) = Q(d|h) = R(d|h) \; . \tag{9}$$

Fine. Now consider those cases where, as it turns out,

$$P(h) = \lambda Q(h) + (1-\lambda) R(h) \; . \tag{10}$$

Letting

$$G(d) = f_{\mathbf{P}}(d) - \lambda f_{\mathbf{Q}}(d) - (1-\lambda) f_{\mathbf{R}}(d) \; , \tag{11}$$

we see immediately that

$$\begin{aligned} G(d) &= \sum_h \Big[P(h)P(d|h) - \lambda Q(h)Q(d|h) - (1-\lambda) R(h)R(d|h) \Big] && (12) \\ &= \sum_h \Big[P(h) - \lambda Q(h) - (1-\lambda) R(h) \Big] P(d|h) && (13) \\ &= 0 && (14) \end{aligned}$$

since

$$f_{\mathbf{P}}(d) = f_{\mathbf{P}}(d) \sum_h P(h|d) \equiv \sum_h P(h) P(d|h) \; , \tag{15}$$

and so on.

And that's it, from this perspective.

$$P(h) = \lambda Q(h) + (1-\lambda) R(h) \qquad \Longrightarrow \qquad f_{\mathbf{P}}(d) = \lambda f_{\mathbf{Q}}(d) + (1-\lambda) f_{\mathbf{R}}(d) \tag{16}$$

when and only when the agents who hold the beliefs **P**, **Q**, and **R** "agree" to the meaning of the d-"clicks" the measuring device will elicit.

Part of me fears that you're going to view all this as little more than rhetoric. It took me five pages of explanation to do what you did in one paragraph. So let me try to reiterate the point of this exercise from my perspective one last time, before I throw in the towel. Here's the point. If you believe that the quantum state is rigidly (or even loosely) connected to reality, then—it seems to me—you will never find a way out of the conundrum of "unreasonableness" associated with "state-vector collapse at a distance." I.e., our community will always be left with a search for the *mechanism* that makes it go. Our community will always be left with the embarrassing questions to do with its clash with Lorentz invariance. And, maybe most importantly, we will be left with the nagging question of why we can't harness this mechanism for more useful purposes.[12] I view

[12] If we can't see angels, why posit them? If we can't see superluminal communication, why posit it? Alternatively, if we do posit angels, the natural question to ask is how might they help save our souls. I would be suspicious of a world where the angels served little purpose other than rounding out a theology and not aiding in our souls' deliverance.

these questions as a distraction from the ultimate goal we all ought to have, of building a more interesting, more fantastic physics.

In connection to the present discussion, I would claim that THE preparation is little more than an anchor for such a rigid connection. It is, even if implicitly, a reversion back to the old issues that led to all the trouble in the first place. Thus I am compelled to find a way to absolutely disconnect the quantum state from the physical world. Yet I am still required to make sense of what it is that we're doing when we practice quantum mechanics. For this, the trail has already been blazed by (radical) Bayesian probability theory, and thus I take that as my cue. Getting at Eq. (16) from within a subjectivist framework is one of those first steps you just have to take ... and then hopefully never have to think about again.

OK, with that, I'm ready to praise you my third and final time in this letter. But I think you'll probably hardly feel it's a praise. Let me tell you another goal of seeing how much of quantum mechanics can actually be run through with complete airtight consistency from the subjectivist standpoint. And that is to pick the theory apart. For, you see, I see no difficulty whatsoever in imagining that any theory can consist of two basic components—the completely subjective and agent-dependent, and the completely objective (or intersubjective if you will) and thus, agent-independent. What is the distinction between these components (I hope you ask)? It is this: Once an agent posits one of the objective components in any particular instance (whether it's "really" there are not), there is no move *within* the theory that will allow him to change that supposition.

Let me give you an example of the latter. Posit a bipartite system with Hilbert spaces $\mathcal{H}_{d_1}$ and $\mathcal{H}_{d_2}$ (with dimensions d_1 and d_2 respectively) and imagine an initial quantum state for that bipartite system. Now, I would say that the quantum state must be a subjective component in the theory because the theory allows me localized measurements on the d_1 system that can change my quantum state for the d_2 system. In contrast, I would say that the number d_2 that was posited for the second system is an objective component in the theory. There is nothing I can do at the d_1 site that will allow me to change the numerical value of d_2. The only way I can change that number is to scrap my initial supposition. Thus, to that extent, **(on first pass)** I have every right to act as if the numbers d_1 and d_2 are potential "elements of reality."[13]

Here's where I really think you sell yourself short by advertising your system as an extension or generalization of classical probability theory (with classical probability theory as a special case that's gotten by deleting one of the axioms). For I would say that your framework of "states" as vectors and "measurements" as applications of Bayes' rule is *classical probability theory*, full stop. Or, I should just say "*probability theory*, full stop"—for, the word "classical" seems to imply that it is a subject somehow within empirical science (rather than "law of thought" that antecedes science). In showing me that even quantum "measurements" can be viewed legitimately as nothing more than applications of Bayes' rule, you have done me a great service. For you demonstrate to me more clearly than ever that the concept of POVM ought to be put onto the subjective side of the shelf when I tear quantum mechanics into its two components. But your other intriguing axioms—like the simplicity and composite-system axioms—which you think give the possibility of generalizing upon classical probability, I would say are nothing of the sort. Instead, I would say they express just the opposite. These axioms seem to me to say something about what we are positing of nature. They express something that is not subjective and is not "law of thought." They in fact form part of the restriction to probability theory that I asked for over and over in my `quant-ph/0106166`. Thus drawing those axioms out explicitly strikes me as great progress.[14]

[13]Please note that I emphasized the qualifier "on first pass." The reason for this emphasis will become clear if you read the letters to Preskill and Wootters that I have pasted into the present letter to you.

[14]And that is why this minor spanking counts as a praise!

Thus to come back to my example of Hilbert-space dimension. As I have already explained, I would say that quantity is a (potentially) objective feature of nature. But now, after understanding your paper better, I would say the same of your composite-system axiom. In fact, that axiom is a way of elucidating the very meaning of the dimension d. As such, it forms part of the "manifold-structure analogy" behind quantum mechanics which I tried very hard to explain to Preskill and Wootters in two further letters I'll paste below. (I hope you will read them along with this note, as I think they greatly elucidate what I was hoping to convey in the previous four paragraphs.)

There. Three praises; take them for what they're worth. They are the three things that really intrigue me about your paper. But I've only started my digestion process. I'll leave a discussion of the confusions until we get some single-malt in front of us next week.

15 March 2002, to Bennett, "Ice Cream and Reciprocity"

Thanks for taking care of Lucien last night. After a night of alternating between sweat and shivers, I seem to be on the road to recovery. A comment and a question.

Comment. I think you said that I should make an overlay for my quantum-axiom slide that says "Give an ice-cream reason, if possible," for each of the axioms. By this, you were indicating that my "Give an information theoretic reason, if possible," is a rather arbitrary thing to be asking. But, I say that it is not arbitrary precisely because the main object of our attention in the theory, the quantum state, specifies *probabilities*. It specifies how we ought to be taking nature into account when we make our mortal decisions. The quantum state does not specify flavors of ice cream. Thus it seems to me like an entirely natural question to ask: If the main object of the theory is of an information theoretic (or decision theoretic or call it what you will) character, then how much of its support structure might also be of the same character?

Question. On the other hand, I rather liked what you said about wanting to base quantum mechanics on the idea that interactions are more symmetric in this theory than in classical physics. In interaction, both observer and observed are changed in the process. Could I get you to write your own version of that in a small paragraph, so that I can have something solid to read and think about?

27 March 2002, to Gea-Banacloche, "Still More Zing!"

Thank you for the long, beautiful note. I loved it, and it makes me so wish that I had had enough time to get to the rest of my talk in Dublin—i.e., the part where I try to give some substance to the word "Zing!" I introduced on one of the early slides. (Recall, "Zing!" was meant to be a place holder for the answer to the question "What is real about a quantum system?")

Gea-Banaclochesm 1: *You seem to have a pretty good idea of how to make most of the postulates (on probability, tensor space structure, and even wavepacket reduction) flow in a more or less natural way from some reasonable information-theoretic ideas, once you are given the basic formal structure of operators and Hilbert spaces. The main question would appear to be, where does this formal structure come from, and what does it actually say about the physical universe?*

That is the main question. And—*in spirit*—I believe our proposed answers appear to be essentially the same:

Gea-Banaclochesm 2: *the basic physical fact at the heart of quantum mechanics is the uncertainty principle, which one could formulate very generally as follows–*

(P1) The nature of quantum mechanical systems is such that, even when we have all the possible information we can have about them, we cannot, in general, predict the outcomes of all the possible experiments we could carry on them. Specifically,

The only issue in my mind is how to carry out the word "specifically" in a way that would satisfy the aesthetic I'm seeking. In particular, I would really like to pin down where the noncommutativity comes from in a way that does not make a priori use of the notions of "compatibility" and "incompatibility." In other words, I'd even like "incompatibility" to be a secondary notion, rather than a primary one. I think it is possible, and my present thinking is that it can be made to come out in a pretty way as a natural consequence of the mismatch between the number of bits that can be reliably stored in a quantum system and the number of measurement outcomes required for an "informationally complete" representation of a quantum state. (That is to say, something along the lines of the mismatch between the numbers N and K in Lucien Hardy's treatment in `quant-ph/0101012` and `quant-ph/0111068`.)

In any case, I have a discussion of "Zing!" in several places in my paper `quant-ph/0106166`. Especially in the closing section. Since then, I've come quite a way toward what I was trying to express above, but you'd have to dig harder for that—it's not exactly published properly yet. The place to look at the moment is in the file titled "Quantum States: What the Hell are They?" posted on my webpage (link below). The upshot is the following (working) statement: Each quantum system can be postulated to have an intrinsic amount of "sensitivity" to our experimental interventions upon it, and that sensitivity can be captured by a single parameter d (call it the dimension). From that, everything about the algebra of observables for a system follows from a basic statement about the very meaning of Bayes' rule in that context.

I'm working hard to get some of this in a proper paper presently.

By the way, I loved your Teilhard de Chardin quote:

> The history of the living world can be summarized as an elaboration of ever more perfect eyes within a cosmos in which there is always something more to be seen.

In my own way, I tried to express something similar in two pieces that I've come to be pretty proud of. Have a look at my letter to John Preskill titled "Psychology 101" starting on page 143 and my letter to Bill Wootters titled "A Wonderful Life" starting on page 149 of the file I mentioned above.

I hope that you yourself made it home safely and comfortably from Madrid, and also that you found your family doing well there. Emma's chicken pox are already clearing up: So maybe I was gone just the right amount of time!

28 March 2002, to Duvenhage, "Quantum-Information Information"

I just want to write you a very short note to tell you how much I enjoyed your paper `quant-ph/0203070`, "The Nature of Information in Quantum Mechanics" and to express how much similarity I think I see between our points of view. In particular, I think you expressed some things so very clearly that I would love to co-opt your phrases!

Here's where I think we agree the most:

1. "A measurement is by definition the reception of information by the observer."

2. "This renders many problems surrounding the measuring process in quantum mechanics no more difficult than in classical physics." And consequently,

3. Your discussion of the Heisenberg cut.

You can find some reflection of these ideas in my own paper quant-ph/0106166, though not put quite so succinctly as in yours. In particular I'm thinking of my discussion on pages 27 and 28 of that paper, following the earlier discussion on page 11. Also, I agree with your point about B's information being "invalidated" in your discussion on page 13 of your paper. Similar ideas make an appearance on pages 39 through 41 of mine. Finally, I also enjoyed your discussion of the linearity of time evolutions. That was the sort of thing, I was trying to describe in my notes of 22, 26, and 27 September 1999 on pages 408, 284, and 285, respectively, of my samizdat quant-ph/0105039.

However, I think I've come a long way since that paper and those notes. In particular, I think I'm no longer really willing to say that "quantum collapse is a noncommutative [generalization of] conditional probability." I think there is a sense in which quantum collapse is *precisely* an application of Bayes' conditionalization rule, modulo only a final redefinition of what the posterior probability is relevant to. What I mean by this in more detail can be found in my letter to Lucien Hardy beginning on page ??? of my collection "Quantum States: What the Hell Are They?" (which can be found on my webpage, link below), especially toward the end of the letter. In fact, I'm presently striving to write that up in an updated version of quant-ph/0106166, and I hope to place it on the server soon.

Anyway, I really want to point out the similarities in our thoughts rather than the contrasts. I think you've done the physics community a good service with your paper. It's very well written, and a lovely piece actually.

29 March 2002, to van Fraassen, "Building with Bayes"

Sorry to take so long to acknowledge your note: I've been running around Ireland with some bad phone connections and probably a few too many pints.

But I'm back now, and I did find your flyer waiting in my mailbox. I tacked it to the wall, but the subject's not likely to attract any of the physicists in my immediate vicinity. I would like to bring a visiting student with me though. Her name is Petra Scudo, and her present email address is scudo@techunix.technion.ac.il. She, as things turn out, did some undergraduate work in Pavia under a guy named Regazzini, who in turn was a student of Bruno de Finetti. While Petra is visiting (for a month and a week), I'm going to have her work on trying to pin down a kind of quantum de Finetti representation theorem for time-evolution maps. I.e., a theorem in answer to the question, "What is an 'unknown' quantum operation?" (In this context, the term "quantum operation" refers to the generalization of unitarity that is common in quantum channel theory — namely, the trace-preserving completely positive linear map. Sometimes people call it "open-system dynamics" but, from the Bayesian perspective, it is little more (nothing more?) than a noncommutative generalization of a conditional probability, connecting as it does the prior (quantum-state) assignment to the posterior (quantum-state) assignment.

Technically, I've made a breakthrough of sorts recently. I now know how to think of quantum collapse as *precisely* an application of Bayes conditionalization (importantly, followed by a redefinition of which measuring instrument the posterior probability assignment is relevant to). I call this a breakthrough because until recently (for instance see page 25 of my NATO paper, quant-ph/0106166), I continued to think of collapse as a noncommutative *analogue* of conditionalization. But now, using a representation of the quantum state similar to the one Hardy harps on, I can see that what is going on is the true-blue thing (i.e., simply Bayes in disguise). At the moment I'm working hard to get this written up in a sensible way ... or at least give my readers a hint of it, until I can do it properly.

By the way, in coming to all this, I've taken a more radical Bayesian turn than I had expected at the outset. I.e., though I started my career in an Ed Jaynes kind of "objective Bayesian" camp, I'm now finding myself in the left of the de Finetti camp and maybe a little beyond that. In case it interests you, I've documented a lot of this transition in a new samizdat which I've placed on my webpage `http://netlib.bell-labs.com/who/cafuchs/`. The title of the file is "Quantum States: What the Hell Are They?" and it contains a lot of new correspondence with Mermin and others along these lines. (As an aside, I've significantly revamped my webpage; it's not so minimalist anymore, and maybe thus a little more attractive.)

While I'm here and I've invested this much into a long letter to you, let me make the thing even longer by tacking on two pieces from the above-said collection. I'll put them immediately below—one is a letter to John Preskill and one a letter to Bill Wootters. Both letters should be self-contained. Anyway, I place them here because I had forwarded them to Henry Folse a while ago, and he wrote me back a rather excited letter saying that I'm starting to explore some philosophical ground not so dissimilar from where Bas van Fraassen has gone. I wish I were in a position to judge the validity of that! But I just haven't read enough yet. (I'm trying to, believe me. But, being a physicist, I've got a lot of material to catch up on.) Anyway, until then, I'll keep my fingers crossed that maybe I can get some reaction directly from the horse's mouth. Are there similarities between our views? And what other pieces of your work should I be reading if there are?

Finally, concerning your seminar with Elga, I'd love to attend! So, please do keep me on the list.

06 April 2002, to Plotnitsky, "The Early Morning"

I laid in bed this morning, reading and thinking about your review of Feyerabend's book. What a pleasant way to wake up. I think you are right: the most important issue is, what is the right balance of construction? It plays a role—I am convinced—but how big of a role?

You flatter me by sending me your newest book! How can you be so productive, while all the rest of us just twiddle by? I will cherish the book and have it read before you know it.

09 April 2002, to Folse, "Thanks"

> **IMPORTANT**: See the correction to the present note in my note "Doing It and Doing It Right," dated 15 April 2002, to Henry Folse.

Folse-ism 3 :

> That is, the quantum state is a compendium of Bayesian "beliefs" or "gambling commitments" and is thus susceptible to the type of analysis James gives in his "Sentiment of Rationality." Our particular choice of a quantum state is something extra that we carry into the world.

Whoa, this paragraph eludes my comprehension. I thought the measurement chose the state of the outcome.

That's a long story. Yes, measurement's can determine states, but my latest greatest realization is that they determine states in a way not so different from the way probability distributions $p(x)$ and $p(y|x)$ determine $p(y)$. $p(x)$ is the classical analog of the initial quantum state. $p(y|x)$ is the classical analog of the projection operator in the collapse rule. $p(y)$ is the classical analog of

the final state. To the extent that in the Bayesian view *all* probabilities—even conditional ones like $p(y|x)$—are completely and utterly subjective, so are the meanings of the outcomes of one's quantum measurement (i.e., so are the basis vectors to which the collapses occur).

10 April 2002, to Bub, "Bitbol-ization"

Bubism 1: *What seems to me closely related to Bitbol's position is your emphasis on replacing the idea of a physical theory (you say 'law') as a mirror image of what is in the world (in what Pauli called a 'detached observer' sense), with an opposing view that takes seriously the fact that our science can only reflect points of view we can construct from within the world. I agree with you that this is the right way to look at quantum mechanics (although in my book I argued for the 'detached observer' view, which I identified with an Einsteinian view). Now I would say that quantum mechanics is not so much a descriptive theory of new sorts of non-classical objects (particles that are also wave-like, or particles with properties that hang together in a non-Boolean way, for example), as a theory of mechanics constrained by certain explicit limits to the process of objectification. (So, as Bohr said, there are no quantum objects. Quantum mechanics is not about how nature is, but about what we can say about nature. In this sense, it's a mechanical theory at the information-theoretic level, unlike classical mechanics – and the claim is that we are stuck at this level just because we are not 'detached observers.') Bitbol talks about the 'blinding closeness' of the world in the title to one of his books – a much more apt image than d'Espagnat's image of a 'veiled reality.'*

Sorry myself for taking so long to get back to you. Bitbol sounds interesting. I've put in a good word for him to Gilles; we'll see what happens. I absolutely love that phrase "blinding closeness."

Bubism 2: *Back to Bitbol and your earlier letter to me with replies to Preskill and Wootters. The essential point there seems to me very close to the neo-Kantian view about quantum mechanics that Bitbol has been developing in several books and articles.*

At the risk of getting into deep waters with a philosopher, I would say the point I'm pushing has a much deeper affinity to the philosophical tradition of pragmatism (James and Dewey's versions in particular) than anything of a Kantian flavor. But I've learned my lesson about saying the words "Copenhagen interpretation" to physicists. (See the story on the bottom of page 70 in the big Samizdat.) And similarly I'm learning a lesson about saying the words "James," "Dewey," and "pragmatism" in front of philosophers. The reactions I've gotten from Timpson, Brown, Donald, and Butterfield! Even when, in other contexts, I got such pleasant reactions from them about the ideas I was talking about ... namely pragmatism without the explicit label. Sometimes a few words are worth a thousand (mental) fences.

11 April 2002, to Caves-Deutsch, "Zing D"

I am finally writing up my contribution to the Växjö conference proceedings, and in it I plan to make a statement about which elements of quantum theory I would be willing to call "ontological" if push came to shove. As Carl knows well, my favorite is the D. That is, for each system, when I ascribe to it a Hilbert space of some dimension D, what I am really doing is ascribing it an integer parameter of some (potentially) ontological significance. I do not let myself, however, assume such a significance for the states *in* the Hilbert space or the operators *on* it.

When I have that discussion, I would like to cite the stuff I saw Ivan present at the ITP meeting. (I saw it via the web.) Do you have paper written on that? If so, what are its coordinates? What I'll say is something like: If you're looking for the magic ingredient that powers quantum computing, it's not going to be first and foremost in the subjective elements of the theory. It's going to be in those things that stand a chance of being objective.

By the way, Ivan. I remember seeing one slide where you talked about the various points of view for just the question above—i.e., what might give quantum computing its power. In it you had one bullet devoted to "information-disturbance tradeoff." That was flattering, but surely it must be such a minority opinion as to not be an opinion at all!! (However, you ought to know what I'm aiming for: that that and D are the same thing, after all. So thanks in advance!)

13 April 2002, to Garisto, "Quantum Locality"

Thanks for sending me your draft; I will try to have a deeper look at it after my present project lets me come up for a bit of air. But, yeah, you're right, I can already tell I'm going to disagree with it. "What is the speed of quantum information?" It doesn't have a speed. It could only have a speed if you endow the state vector with an objectivity it does not have. QUANTUM STATES DO NOT EXIST (in and of themselves). They merely express the gambling commitments one is willing to make when one encounters a quantum system. But that does not leave the world empty; it just means that the quantum state is not part of its substance. Why is it that I should choose my gambling commitments to be in accordance with the structure of quantum mechanics? When we can answer that, we will have learned something clean and simple about the substance of the world. But until then, encumbering the world with an idea—nonlocality—that is clearly wrong-headed (without a heck of a lot of contrivance to try to make it go) will only distract us from the straight course to that great goal. What is it that makes quantum mechanics go? It is something deeper and far more interesting than the quantum state—that much I firmly believe.

I'm glad to hear that things are going well with you cancer-wise. Every time I think of you, I cross my fingers mentally for your health and happiness.

15 April 2002, to Folse, "Doing It and Doing It Right"

I was just looking my last note to you over again, and I was appalled at what I had written. Namely, I really botched my description of the classical analog to collapse. How, I could do that, I don't know. And I am ashamed of myself.

Here's the proper way to say it, and in fact the way I am just writing it up for a paper:

> In Section 8, "**What Else is Information?**," I argue that, to the extent that a quantum state is a subjective quantity, so must be the assignment of a state-change rule $\rho \to \rho_d$ for describing what happens to an initial quantum state upon the completion of a measurement—generally some POVM—whose outcome is d. In fact, the levels of subjectivity for the state and the state-change rule must be precisely the same for consistency's sake. To draw an analogy to Bayesian probability theory, the initial state ρ plays the role of an a priori probability distribution $p(h)$ for some hypothesis, the final state ρ_d plays the role of a posterior probability distribution $p(h|d)$, and the state-change rule $\rho \to \rho_d$ plays the role of the "statistical model" $p(d|h)$ enacting the transition $p(h) \to p(h|d)$. To the extent that *all* Bayesian probabilities are subjective—even the probabilities $p(d|h)$ of a statistical model—so is the mapping $\rho \to \rho_d$. Specializing to

the case that no information is gathered, one finds that the trace-preserving completely positive maps that describe quantum time-evolution are themselves nothing more than subjective judgments.

16 April 2002, to Bennett-Brun-Caves-Grangier-Mermin, "My Own Homeopathy"

I thought of all of you with a smile as I was writing the footnote below for an upcoming paper. I'll send it to you now. Maybe I really do practice homeopathy. (Wait till you see the undiluted version.)

> In the previous version of this paper, `quant-ph/0106166`, I variously called quantum states "information" and "states of knowledge" and did not emphasize so much the "radical" Bayesian idea that the probability one *ascribes* to a phenomenon amounts to *nothing* more than the gambling commitments one is willing to make with regards to that phenomenon. To the "radical" Bayesian, probabilities are subjective all the way to the bone. In this paper, I start the long process of trying to turn my earlier de-emphasis around (even though it is somewhat dangerous to attempt this in a manuscript that is little more than a modification of an already completed paper). In particular, because of the objective overtones of the word "knowledge"—i.e., that a particular piece of knowledge is either "right" or "wrong"—I try to steer clear from the term as much as possible in the present version. The conception working in the background of this paper is that there is simply no such thing as a "right and true" quantum state. In all cases, a quantum state is specifically and only a mathematical symbol for capturing a set of beliefs or gambling commitments. Thus I variously call quantum states "beliefs," "states of belief," "information" (though, by this I mean "information" in a more subjective sense than is becoming common), "judgments," "opinions," and "gambling commitments." Believe me, I already understand full well the number of jaws that are going to drop by the adoption of this terminology. However, if the reader finds that this gives him a sense of butterflies in the stomach—or fears that I am or will become a solipsist[1] or a crystal-toting New Age practitioner of homeopathic medicine[2]—I hope he will keep in mind that this attempt to be absolutely frank about the subjectivity of **some** of the terms in quantum theory is part of a larger program to delimit the terms that actually can be interpreted as objective in a fruitful way.

1. P. Grangier, private communication, 2001.
2. C. H. Bennett, T. A. Brun, C. M. Caves, and N. D. Mermin, various vibes, 2001.

23 April 2002, to Comer, "Music in the Musician"

Thanks for the thoughtful note. I found myself thinking about it on and off all last night, in both my periods of wake and sleep. I think you expressed the issues to do with chemistry versus consciousness especially vividly.

I think we just have to get rid of this imagery that we are "made" of atoms. Or none of us are ever going to make any progress in our emotional lives OR our physical understanding. By my present thinking, a much better imagery is this. Take me and an old log: we both float in water.

That is to say, we have that much in common. But there are a heck of a lot more things that we do not have in common. For any two entities, we can always find some characteristics they have in common, if WE are willing to ignore all the ways in which they are distinct. And that, I think, is the story of atoms. The atomic picture has something to do with what we all have in common. (Or, maybe more potently, it has something to do with what is common in the PART of existence that we have chosen not to ignore for the moment.) But to see the atomic picture shine through, we have to dim down all the things that are unique in us. Who said the particular shape of that rock is not important? Who said the pain you are feeling is only epiphenomena?

Such a picture of what physics and chemistry is about is every bit as consistent as the worldview Steve Weinberg, say, would have us believe. And I would say that it is more so; for it gives us a power and a hope for control in our lives that his can't even imagine.

Let me do two things for you. First, I'll paste in two old emails, that have to do with your music-on-the-page versus music-in-the-musician imagery (which I think it is so apt and so beautiful). Mostly I'm pasting them because your note caused me to go back and read them this morning. And I'm just reconfirming that I'm on the same wavelength.

But then I want to quote William James. (That will come in a little later note.) I know you're not much in the mood to read any philosophy right now. But if you read the RIGHT stuff, I cannot see how it cannot help. My side of the conversations with you, in any case, is just a poor reflection of what William James already said with such flare.

17 December 1997, to Comer, "It's a Wonderful Life"

Good holidays to you. This morning, as I was driving to work, it dawned on me that roughly this day 10 years ago, I was conferred my degrees at the University of Texas. Time does fly.

It made me think of a little anecdote about John Wheeler that I heard from John Preskill a few days ago. In 1972 he had Wheeler for his freshman classical mechanics course at Princeton. One day Wheeler had each student write all the equations of physics s/he knew on a single sheet of paper. He gathered the papers up and placed them all side-by-side on the stage at the front of the classroom. Finally, he looked out at the students and said, "These pages likely contain all the fundamental equations we know of physics. They encapsulate all that's known of the world." Then he looked at the papers and said, "Now fly!" Nothing happened. He looked out at the audience, then at the papers, raised his hands high, and commanded, "Fly!" Everyone was silent, thinking this guy had gone off his rocker. Wheeler said, "You see, these equations can't fly. But our universe flies. We're still missing the single, simple ingredient that makes it all fly."

Merry Christmas.

02 December 1997, to Preskill, "Flying Equations"

I couldn't help but think of the anecdote about John Wheeler's (non-)flying equations you told the other day when I came across the following little passage (presumably Biblical in origin):

> "I forbade any simulacrum in the temples because the divinity that breathes life into nature cannot be represented."

23 April 2002, to Comer, "Installment 1"

From *Pragmatism*, pages 30–32:

> And do not tell me that to show the shallowness of rationalist philosophizing I have had to go back to a shallow wigpated age. The optimism of present-day rationalism sounds just as shallow to the fact-loving mind. The actual universe is a thing wide open, but rationalism makes systems, and systems must be closed. For men in practical life perfection is something far off and still in process of achievement. This for rationalism is but the illusion of the finite and relative: the absolute ground of things is a perfection eternally complete.
>
> I find a fine example of revolt against the airy and shallow optimism of current religious philosophy in a publication of that valiant anarchistic writer Morrison I. Swift. Mr. Swift's anarchism goes a little farther than mine does, but I confess that I sympathize a good deal, and some of you, I know, will sympathize heartily with his dissatisfaction with the idealistic optimisms now in vogue. He begins his pamphlet on 'Human Submission' with a series of city reporter's items from newspapers (suicides, deaths from starvation, and the like) as specimens of our civilized regime. For instance:
>
> "After trudging through the snow from one end of the city to the other in the vain hope of securing employment, and with his wife and six children without food and ordered to leave their home in an upper east-side tenement-house because of non-payment of rent, John Corcoran, a clerk, to-day ended his life by drinking carbolic acid. Corcoran lost his position three weeks ago through illness, and during the period of idleness his scanty savings disappeared. Yesterday he obtained work with a gang of city snow-shovelers, but he was too weak from illness, and was forced to quit after an hour's trial with the shovel. Then the weary task of looking for employment was again resumed. Thoroughly discouraged, Corcoran returned to his home last night to find his wife and children without food and the notice of dispossession on the door. On the following morning he drank the poison.
>
> "The records of many more such cases lie before me [Mr. Swift goes on]; an encyclopedia might easily be filled with their kind. These few I cite as an interpretation of the Universe. 'We are aware of the presence of God in his world,' says a writer in a recent English review. (The very presence of ill in the temporal order is the condition of the perfection of the eternal order, writes Professor Royce (The World and the Individual, II, 385).] 'The Absolute is the richer for every discord and for all the diversity which it embraces,' says F. H. Bradley (Appearance and Reality, 204). He means that these slain men make the universe richer, and that is philosophy. But while Professors Royce and Bradley and a whole host of guileless thoroughfed thinkers are unveiling Reality and the Absolute and explaining away evil and pain, this is the condition of the only beings known to us anywhere in the universe with a developed consciousness of what the universe is. What these people experience *is* Reality. It gives us an absolute phase of the universe. It is the personal experience of those best qualified in our circle of knowledge to *have* experience, to tell us *what is*. Now what does *thinking about* the experience of these persons come to, compared to directly and personally feeling it as they feel it? The philosophers are dealing in shades, while those who live and feel know truth. And the mind of mankind—not yet the mind of philosophers and of the proprietary class—but of the great mass of the silently thinking men and feeling men, is coming to this view. They are judging the universe as they have hitherto permitted the

hierophants of religion and learning to judge *them*. ...

"This Cleveland workingman, killing his children and himself [another of the cited cases] is one of the elemental stupendous facts of this modern world and of this universe. It cannot be glozed over or minimized away by all the treatises on God, and Love, and Being, helplessly existing in their monumental vacuity. This is one of the simple irreducible elements of this world's life, after millions of years of opportunity and twenty centuries of Christ. It is in the mental world what atoms or sub-atoms are in the physical, primary, indestructible. And what it blazons to man is the imposture of all philosophy which does not see in such events the consummate factor of all conscious experience. These facts invincibly prove religion a nullity. Man will not give religion two thousand centuries or twenty centuries more to try itself and waste human time. Its time is up; its probation is ended; its own record ends it. Mankind has not aeons and eternities to spare for trying out discredited systems."

25 April 2002, to Bennett, "Short Thoughtful Reply"

Thanks for the picture of the skunk cabbage. I've always wanted one. I'm sorry I wasn't able to reply to you earlier, but with all the students visiting, etc., I've had a gazillion things going on at once.

Let me give you a very short reply, for what it's worth.

Bennettism 2: *My main wonder about your beliefs is why do you find it so important to emphasize the subjectivity of quantum states, ... What difference does it make in any case?*

I am just trying to do what scientists always try to do: understand how things "hang together." I.e., build a (satisfying) picture of the world that has nothing to do with my personal qualities. It just so happens, that my favorite problem happens to be different from your favorite problem.

Bennettism 3: *Do you think Katie really exists, or is she just a mathematical symbol for a set of bets you would [be] willing to make?*

Of course I think there is a sense in which Katie exists (i.e., some large remnant of Katie as she is now would be here even if I were killed tomorrow). I just make a distinction between all the stories I might write about her in my samizdats and whatever it is that she *is* in herself. What I don't understand about you is why you find that such a foreign concept.

In particular, I don't think I could ever write a sentence like this:

Bennettism 4: *My main wonder about your beliefs is why do you find it so important to emphasize the subjectivity of quantum states, but not other kinds of information, such as the dinner you just ate, your shoes, or other people?*

I don't think I've ever thought of the dinner I just ate as the information I just ate. Presumably there is something substantial to broccoli independently of my subjective judgments about how it tastes. Information, as motivated by Shannon, has something to do with the concept of surprise. If I believe strongly that broccoli tastes bad, then I will be surprised if I find that I actually like it. In that sense I will find that I have gained a lot of information when my subjective judgment makes a transition from its old value (Yuck!) to its new value (Mmm!). But that has nothing to do with broccoli as it is completely independently of me.

Even if I thought of a quantum state as an objective rather than a subjective quantity, I still don't think that I could ever talk as you did above. There is a difference between "systems" and "properties." And there is something in your language that seems to blur the distinction.

To make this concrete, take a classical description of a pendulum's motion in terms of phase space. I would never call the phase space point (x, p) "the particle." The particle is what *carries* the property (x, p). Within classical physics, both the particle and its property might as well be assumed to exist even when there are no physicists about. But let me ask the same thing about the Liouville distribution for a *single* instance of the particle. Can the single particle be said to "carry" a Liouville distribution in the same way it "carries" the coordinates (x, p)? I think you would be hard pressed to say that it does. For if I were to delete the physicist who is ignorant of the phase-space point the particle actually possesses, then I would delete the Liouville distribution. But I would not delete the value (x, p)—whichever one it is—that the particle can be safely assumed to have.

And that is all I am striving for in quantum physics. To figure out what properties we might safely ascribe quantum systems even when there are no physicists about. I think there are awfully good reasons for thinking that "the" quantum state is not such an "objective" thing. And thus the quantum state carries more of an analogy to the Liouville distribution than to the phase space point. However, that does not mean that I think our beloved theory gives us no hint of what the properties are that I can safely treat as objective. It is just that, among them, I do not see the quantum state.

Bennettism 5: *When you say a quantum state is just a set of bets you would be willing to make, what is the ontological status of you the bettor? Are you just a collection of bets some other bettor would be willing to make?*

Don't blur the distinction between the system and the state!

There, that's my short reply. As I say, I wish I could have replied to you earlier, but I had so many things tying me up. I wonder if I could make a birthday wish of you? Since my birthday was the 21st, would you give me this much of a present? Just read the parts of the paper I posted for `quant-ph` tomorrow, to do with Preskill and Wootters. I'll send it to you shortly. It's not long reading, and it's not hard reading. (Certainly no harder than a New Yorker.) And just give me two binary digits of satisfaction: After reading those passages, would you 1) say that you still do not understand my views of Katie, and 2) does it still look like a tower of turtles to you?

With enduring friendship (and a picture of skunk cabbage hanging on my wall),

03 May 2002, to van Fraassen, "Accepting Quantum Mechanics — The Short of It"

I'm feeling horribly guilty because I promised you a note over two weeks ago and everything—just everything—has conspired against me getting it constructed. I'm sorry. But still I want to write you something before I meet you next week. Let me try to do what I can in the next hour and then call it quits until Monday: The details of what I am about to say will be in a paper that I plan to have finished and can bring with me then.

van Fraassenism 1: *Now when it comes to theories that give us probabilities, whether absolute or conditional, I'll agree with scientific realists that literally read they say that there are objective probabilities in nature. But accepting such a theory does not involve believing that. Rather it involves appointing the theory as an 'expert' for guidance of our subjective probabilities concerning*

observable events. The metaphor of 'expert' is cashed out (as by Haim Gaifman) as follows. Suppose that I appoint Peter as my expert on snuffboxes. That means for my subjective probability P and Peter's subjective probability q the constraint:

$$P\Big(A \,|\, q(A) = x\Big) = x \tag{17}$$

with generalizations of this to intervals, odds, conditional probabilities, for statements A that are about snuff boxes.

Thus the issue of whether there are objective probabilities in nature or whether to believe in them is finessed: there are only the theory's probabilistic pronouncements *accepted as input and my own subjective probabilities.*

That is clearly not how you are approaching it overall. But perhaps there are connections? I'd like to know how the QM probabilities are fed into your subjective probability as a whole – I wonder if it will not be similar. After all, even if a quantum state is read as a compendium of probabilities, and you say something like "this material is in quantum state such and such", your own subjective probability function has a domain much larger than facts pertaining to this material.

Yes, you are right, I don't like (ultimate) experts, and I don't think quantum mechanics has any more need of them than weather forecasting, say. In fact, I think any attempt to hold on to objective probabilities—even in the finessed form that you talk about, where there is a higher authority in whose judgment we place our faith—will only get in the way of our finding the deeper heart of quantum mechanics. Thus, I hold fast to the idea that there is no right and true quantum state EVER—just as de Finetti held fast to the idea that there is no right and true probability distribution EVER—and I take it as the very definition of my foundations program to see what is left behind. From my view, the theory does not pronounce probabilities; it only pronounces what we ought to be doing with them once we have set them just as subjectively as the next guy. (And to make the issue as pointed as possible, I even mean this for pure quantum states.)

What does it mean to accept quantum mechanics then? The imagery I am starting to build looks something like this.

If one generalizes the notion of quantum measurement to the one that has become essential in quantum information theory—namely to the positive operator valued measure (POVM)—then, for each quantum system, one can contemplate a *single*, fiducial quantum measurement for which the probabilities of its outcomes completely specify the system's quantum state. That is to say, whenever I write down a subjective probability distribution $p(h)$, for the outcomes h of such a fiducial measurement, I completely specify a quantum state ρ (pure or mixed). Imagine now that that fiducial measurement device is tucked away safely in some vault in Paris at the International Bureau of Weights and Measures.

A quantum state can be viewed as very literally nothing more than my subjective judgment for what would happen if I were to ever bring my quantum system up to that standard measurement device. What now can one say about a real-world measurement device, like one that we might have here at Bell Labs? Well, bringing my quantum system up to it will generally evoke a click of some sort that I might label d. Using all that I believe of the device, all that I might believe of Lucent's technical prowess, etc., etc., I would be a bad subjectivist if I didn't allow the click to update my beliefs about what would happen if I were to approach the fiducial device. Thus I end up with some updated probability distribution $p(h|d)$. (I.e., some updated quantum state.)

In what sense is this subjectivism connected to Bayesianism? This in part is what this paper of mine is about that I'll give you a copy of Monday. One can show that the usual story of quantum collapse can be viewed as the conjunction of two things: 1) Bayesian conditionalization, and 2) a

final rotation of the axes of the probability simplex for the fiducial measurement. That is to say, quantum collapse in this description is only a pretty damned mild generalization of the Reverend Bayes.

But, again, what does it mean to accept quantum mechanics? It is this. If one studies the properties of these kinds of fiducial measurements, one finds that for no initial quantum state (in the usual Hilbert space picture) and no outcome is it ever the case that $p(h) = 1$. That is, when one accepts quantum mechanics, one eschews certainty for the outcomes of a fiducial quantum measurement. In fact, the set of allowed distributions $p(h)$ forms a convex set that is strictly contained within the probability simplex (i.e., the set of all imaginable probability distributions over an appropriate number of outcomes).

Thus, accepting quantum mechanics is not accepting the existence of an expert, but—in large part—accepting the two ingredients above:

1. voluntarily accepting a restricted range for one's beliefs $p(h)$
2. accepting a slightly modified form of Bayesian conditionalization for updating one's beliefs (i.e., standard Bayes + rotation)

The NEED for 1) a restricted range and 2) a minor modification of Bayes, is where I say we should be looking if we want to be looking for the "meaning" of quantum mechanics. What is it about the world as we view it that compels us to accept those two ingredients? That I see as the important question. And the ever more convoluted moves I see from some of our friends who want to hold on to a nonsubjectivist view of the quantum state, I see as a waste of good brain power.

That is to say, I agree with you in that, "Be a realist if you want to be." But I add to it with respect to the interpretation of quantum mechanics, "Don't do it for those parts of the theory where it is not productive to do so." If you're looking for a little realism in quantum theory, fine, but then look for it in a more clever place than in the state vector.

Anyway, that's my present take.

And I lied: that took me an hour and 35 minutes. I hope it's a little clearer at least for the extra time. See you Monday!

PS. You wrote:

van Fraassenism 2: *The two articles of yours that we took up in our discussion group in the fall were clearly only the beginning, and you have now taken the program much farther.*

Can you tell me which two articles you're referring to? That would give me a clearer vision of which views I've changed since your reading and which views I need to be careful not to let be propagated in your mind.

08 May 2002, to Wiseman, "The QMP"

Thanks for the letter, and particularly, thanks for skimming the paper. I chose a provocative title for the paper in an attempt to draw in the crowd, but the content of the two letters to Preskill and Wootters was pretty serious for me. What I'm looking for, in particular, is a way to make sense of science from a point of view that says at the same time, "Don't even think the terms in your theory ARE or CAN BE a reflection of reality. For the universe is big, and your head is small. And maybe reality is not static and unchanging—and thus describable in any finite terms, like in terms of GR and QM—anyway."

That's a sweeping statement, but I think it may be one of the two great lessons of quantum mechanics. (I'll keep my opinions secret on the other great lesson.) Anyway in particular, the

way I see it, we as a community should be working as hard as we can to STOP trying to see the wavefunction as anything of a reflection of nature. Wavefunctions live in our heads. They live and die in our heads. And when they do their changing, they do that in our heads too. That's the point of view I'm trying to run through with as much consistency as possible. I.e., how can I hold to it, and still have quantum mechanical practice be what we all think it is? How can I hold to it, and not have everything that we say of nature be just a dream?

I'm sorry I've delayed so long in replying to you, but I wanted to do it right, and that required that I finish up a paper so that I could point you toward it. The paper's finished now, and I posted it on `quant-ph` this morning; it'll appear tomorrow. But if you wake up before `quant-ph` does, you can also get it at my webpage (link below). The title is, "Quantum Mechanics as Quantum Information (and only a little more)." The parts that are particularly relevant to our discussion are Section 4.2 and all of Section 6.

Wisemanism 1: *He implies that the same problems apply in classical theory. I disagree. The problem is that noncommutative algebra does not apply to the information in our brains. ... I'm not saying that it is impossible to do that, but it is an extremely difficult problem which he dismisses.*

Let me give an analogy that I think is apt. I would say Bohr's great genius in developing his model of the hydrogen atom lay solely in one little move: In dismissing the research program his predecessors laid out before him. That is, in dismissing the idea that the atom's spectrum required a mechanical explanation. Beyond that, the rest of his work was just details. Likewise, I think it is the case here.

I think that wavefunction collapse simply calls out for no explanation at all. And that is because I see it as nothing but a variety (and, in a way, an extension) of Bayesian conditionalization. To that extent, the problem—or, by my view, the nonproblem—was already there long before quantum mechanics ever showed up on the scene. Imagine that physical theory really were like the Newtonians wanted it to be: nice and deterministic, Laplace's demon and all that. Now consider a weatherman immersed in the world. He wakes up in the morning, thinks about all the weather readings he has taken the previous days, complains about the fact that his computer can't do as much number-crunching as he would like it to do, scratches his butt, decides whether he is feeling optimistic or pessimistic (that, of course, might depend upon how his girlfriend has been treating him lately), and so on and so on. He churns that all into a big mental pot, and finally comes up with a set of numbers $p(h)$ to describe his beliefs about what the weather will be doing tomorrow. In fact, at the same time, he'll probably come up with a whole set of numbers $p(h, d)$ to describe not only the weather tomorrow, but also the weather today.

Now, suppose he goes and has a look at the weather today and finds that it is d. Then, using Bayes' rule, he will update his belief for tomorrow from $p(h)$ to $p(h|d)$. It's a gut-wrenching, horribly discontinuous transition. But does it call out for an explanation? And if it does, does it call out for an explanation that has anything to do with the system the weatherman is modeling, i.e., the earth's atmosphere? For after all, in this world, when the weatherman looks out the window, one value of d was true and always remains true; one value of h was true and always remains true. The only discontinuous change is in his beliefs ... and those beliefs are presumably a property of his head.

So, yes, I think Duvenhage is absolutely right on this point. It's a point I've been trying to make for years—it's what Asher and I were up to in our *Physics Today* article—but of course sometimes it takes a long time to get the expression right. And I think Duvenhage did a particularly nice job on this particular score.

Now you and I both know that there's something more going on in quantum mechanics than there is for the weatherman. It's just that I don't think that "something more" is localized in the

issue of "collapse." As far as the discontinuous change of belief goes—if you ask me—that happens both quantumly and classically, and in fact has nothing to do with any particular scientific theory. It is pre-science; it is simply a part of one's living and changing his beliefs in response to stimuli from the world outside himself.

By my present thought, part of the "something more" that goes on in quantum mechanics (as opposed to simple Bayesianism) is that we no longer have the right to assume that the things we do in our laboratories to change our beliefs (i.e., quantum measurements) leave the world unscathed in the process. I.e., we ought to be taking into account that when the world stimulates us, we stimulate it back in the process. But how do we take it into account in our description? My answer is that there's only one place to put it in a formalism, and that is in a further change in BELIEF up and beyond that dictated by Bayes' rule.

Thus, quantum collapse deviates from Bayes' rule, but not because it has anything to with something going on outside our heads. It deviates from Bayes' rule because the subject matter we are talking about when we are doing quantum mechanics (i.e., quantum systems) have an implicit action-reaction principle that Bayes, Cox, Ramsey, and de Finetti overlooked when they first worked out the calculus of belief change—i.e., Bayes' rule.

Now, you say things like,

Wisemanism 2: *The problem is that noncommutative algebra does not apply to the information in our brains. ... If we could believe that non-commutative algebra did apply to our knowledge then of course that would solve the quantum measurement problem. But how can we reconcile that with our experience?*

but, by my present view, using words like that is a red herring. (And Duvenhage is much more guilty of it than you.) The way I would say it now—please read the sections in the new paper I told you about—is that any kind of noncommutativity in quantum mechanics is just an artifact of a certain representation. The usual representation is a useful one to be sure—I could hardly calculate anything without it—but for the present issue I think it detracts from the clearer understanding we can hope to obtain if we'll just suppress it. (The "it" meaning the usual representation.) In fact, I think we're only going to get that understanding by exploring quantum mechanics as 1) a restriction on the space of probabilities (the probability simplex), with 2) a conditionalization rule that goes just a bit beyond Bayesian conditionalization.

And all this causes me to reject it when you say,

Wisemanism 3: *The QMP is to find (i) a cut, and (ii) a way to bridge the cut, between the quantum systems that do have non-commutative information (in his terminology) and classical systems that do not.*

That is not the quantum measurement problem for me. People, experimentalists, scientific agents, have information or lack information when they are trying to talk about something outside themselves. Systems are just systems, and when they are treated as such, I would say the concept of information has nothing to do with them. Information is something I have or lack about a system. For instance, when I am speaking about you—thinking of you as a physical system—I lack quite a bit of information in the sense that your behavior could surprise me at any moment. But that means nothing about any kind of "information" intrinsic to you.

So, when you say what you said above IS THE QUANTUM MEASUREMENT PROBLEM, I would say that that is a problem that comes from trying to think of the quantum state, or the algebra of observables, etc., as literal properties of the things you are describing ... and failing to recognize that the quantum state is a property of the observer and not the system. The quantum

state is the full compendium of gambling commitments you would be willing to make about what the system will cause a measuring device to do.

What I myself see as the quantum measurement problem is to give compelling reasons for the two items I listed above. By compelling, I mean in the terms I lay out in the paper. When we can finally do that, then we will finally understand what properties we are really assuming for the "reality" of a quantum system.

Anyway, I hope that helps explain my position. I'll attach another piece of correspondence below with Bas van Fraassen that takes a different tack on some of the same issues. Maybe that'll help supplement the paper. I hope you'll get a chance to take a look at it. By the way, I quote one of your papers in Footnote 33.

Thanks for giving me the opportunity for trying to say these things a little more clearly. And certainly feel free to question or comment on anything that still doesn't make sense. (Anything that doesn't kill me might just make me stronger!) Hey, I'm coming to Brisbane May 25 to June 17. Do you think we'll get chance to talk sometime during then?

14 May 2002, to Schack, "No BC's Role"

Thanks for the flurry of letters. They provoke a lot of thoughts in me and I am grateful.

I'm not sure how I'm going to reply to you, maybe just randomly (as the details occur to me).

Schackcosm 60 : *In my discussions with classical cryptographers, I am often forced to concede that QKD is really quite limited in scope. You either assume an unjammable channel (which I believe cannot be assumed if you use the internet for communicating) or you use classical authentication, which means you share some initial key. Given these limitations, I am not convinced that QKD deserves such an exalted status as suggested by the Bub quote above.*

Has anybody taken the Brassard/Fuchs speculation any further?

Good point, of course. But I never meant for those two "axioms" to be read so literally. Perhaps my best expression of the idea is captured in a letter to Jeff Bub, 10 December 2000. It starts at page 100 of the old samizdat. The main point is that I see information-disturbance as the key idea, along with the commitment that information can never be completely locked away. Read those passages; I think they'll clear some of this up.

That said, precise versions of the no-bit-commitment ideas are coming to the top of my head again for their foundational value. Namely, as part of the extra assumptions that might get us to proper density operators in the "Wither Entanglement?" entanglement section (rather than simply linear operators). The thing that really powers the no-bit-commitment theorem is the Hughston-Jozsa-Wootters result that localized measurements on one system of two (described by a bipartite pure state) can "induce" any decomposition one wants for the other system's density operator. It turns out that the "pure states" of those other wacky operators in my bipartite-Gleason construction don't necessarily have this property. So, it looks like an assumption of such a nature might get me a little closer to the goal. (Though—even if I found it completely acceptable, and I'm not sure I do—it still wouldn't get me all the way to the goal.) But all of this is a long story, and maybe it'd be easier to talk about at a chalkboard.

14 May 2002, to Schack, "More Toenails"

Schackcosm 61 : *Why are you so harsh on entanglement? In my quantum information lectures, I postulated the tensor product structure, as a natural formalism to deal with local operations on*

several particles. Entanglement is then derived, an unexpected consequence of the tensor product structure. You make this line of reasoning more compelling by showing that there is really no alternative to the tensor product structure. Far from being withered, entanglement emerges invigorated from your analysis.

The main point is that entanglement can be thought of as secondary to the structure of quantum observables on localized systems. To that extent, one realizes that one can focus on the structure of simple observables in one's foundational efforts and forget about entanglement. In other words, it seems to me entanglement is not, as Schrödinger said, "the characteristic trait of quantum mechanics, the one that enforces its entire departure from classical lines of thought." It is derived and secondary.

Now what is the structure of observables? The thing I try to argue in the paper is that a measurement is *anything* that gives rise to *any* convex decomposition of a one's original density operator. In that sense, measurement is nothing more than an arbitrary application of Bayes' rule.

Entanglement thus arises from the more basic idea of Bayesian conditionalization in conjunction with the idea that the allowed probabilities for a standard quantum measurement device do not explore the whole probability simplex.

Let me put two notes below that might clarify where I'm trying to go with this. I'd like to think that they add nothing beyond the paper, but it has been my experience that I can just never say enough. (I know what you're thinking: "You've got it backwards. If you'd just say it all in fewer words! That's what you really need.") [See notes to van Fraassen and Wiseman, dated 3 May 2002 and 8 May 2002, respectively.]

14 May 2002, to Schack, "Emphasis De-emphasis"

Schackcosm 62 : *I am not sure I like your emphasis: You describe what I think is the exciting part as a "further readjustment of the posterior state".*

Yeah, I think I agree with that. There's no doubt that that's the most exciting part for me. That extra adjustment strikes me as capturing our beliefs about how we are stimulating the system (rather than how it is stimulating us). And to that extent, I would like to emphasize it. However, I guess I chose the tack I did because I wanted people to stop thinking of quantum collapse as something so different from "mere conditionalization."

Schackcosm 63 : *Personally I like the concept of a compendium of probabilities better than your bureau, but it is very interesting to see how far one can get with the Bayes rule alone.*

Again, here, it is a point of emphasis. The message I am trying to get across is that the structure of observables comes from Bayes' rule. They (measurements) are not defined independently of it. The SQMD struck me as an effective way to badger that point. "Compendium of probabilities" really stands for "compendium of ways of applying Bayes' rule." That is, I think it builds a stronger case for the idea that it is Bayes' rule all the way down when it comes to quantum measurement. The theory gives us no overt means to identify the objective thing that goes on behind a quantum measurement outcome. The only grounding we have is to *declare* a point somewhere off in the distance for which we will do no further updating. And that is the role the SQMD plays.

14 May 2002, to Schack, "Deletions and Their Obverses"

Schackcosm 64: *I think that taking William James and Darwinism seriously means to acknowledge that quantum mechanics is most likely to be superseded one day by a theory with even more cash value. This new theory may not have any of the features that we regard as the core of quantum mechanics. How then is Darwinism or pragmatism going to tell us anything about the foundations of quantum mechanics?*

You're right, I think you are definitely missing what I was hoping to express. But that probably just means I didn't express it so effectively.

No, I did not mean that Darwinism and pragmatism tell us something about quantum foundations. Just the opposite. I see quantum mechanics as giving us a great hint that there is still something deeper. Quantum mechanics is the first rip in the old fabric that told us our place in the world is a nullity. That is to say, I think quantum mechanics only gives us even better arguments for pragmatism.

Schackcosm 65: *It looks to me as if your desire to find the objective core of quantum mechanics is against the spirit of both Darwinism and pragmatism.*

I think the core of the theory—I don't think I ever used the words "objective core"—is just our best guess of what we cannot change with our present level of skills. It is our best attempt to imagine what it would be like if we were not here. Obversely, when we have ferreted the core out, we will have a quantitative indication of how much of the world we can hope to control (given our present skills, present evolutionary level, etc.)

My desire for delimiting the core is expressed particularly in the passages below (taken from the paper). Maybe this makes no sense without A) watching the movie *It's a Wonderful Life*, and B) reading the "Sentiment of Rationality." Have you done either of these things? Probably not.

Your point certainly makes it clear to me that my paper is not self-contained! (The movie is an American classic and I think a large fraction of Americans have probably seen it; but James is another story.)

14 May 2002, to Schack, "Imaginations"

Continuing again ...

Schackcosm 66: *Actually, I found the part where you mention the selection of traits for your daughter outright disturbing, without any compelling connection to the discussion on quantum mechanics.*

Yes, me too. Because read in the wrong way—and maybe there was no other way to read it—it surely evoked images of Nietzsche (at the least) and Nazi Germany (at the worst).

But on the other hand, I don't know how to draw a meaningful distinction between our tools and ourselves. (Remember my point about the prosthetic hand in the paper?) Should we stop pursuing medical research because it goes against the grain of nature? Should we stop pursuing genetic techniques for controlling AIDS? Should we stop falling in love based on an attraction to our partners' complementary (positive) traits to our own? Traits that we would like to see (even if subliminally) appear in our children?

I guess I say no, no, no. Instead we need to cultivate at the same time a respect for everything in its time. Children of Down's Syndrome, say, deserve respect, not deletion. The Nazi experiment

with nature was an atrocity. But, at the same time, I would say we cannot stop pursuing progress in genetic engineering.

The point you make is a deep one. I have no solution. Only an intuition of fear and promise, both at the same time.

If you get a chance, read Richard Rorty's book *Philosophy and Social Hope*. It's all about fear and promise and quite easy to read.

14 May 2002, to Schack, "More Still"

Schackcosm 67: *Everything you write suggests that you want more than just the theory with the highest cash value: you are looking for some form of absolute truth, something that transcends looking for the theory that makes the most accurate predictions.*

Yes, I guess so. I am looking for an indication that the world can be moved. I am looking for an indication that the only law in physics is the law that there are no laws. I am looking for an indication that the world is still writable.

I think if you want to call those things the pursuit of an absolute truth, you can. But they're only absolute in a pretty negative sense.

See, I told you you provoked a lot of thoughts in me.

14 May 2002, to Bennett and others, "Qubit and Teleportation Are Words"

I doubt I will be of any use in constructing a short dictionary definition for the word teleportation, but let me try to explain my difficulty with the word "property" with regards to both quantum states and quantum entanglement. I do this with a little trepidation, but on the other hand, you're the ones who brought me into this discussion and I feel I ought to say something.

The trouble I have with the word "property" has to do with one of the main points Charlie brought up in his talk this weekend. What instantaneously and physically changes about Bob's system when Alice performs a measurement on hers? Charlie told us *nothing* and I agree with that. But then, I look in the *American Heritage Dictionary* and find:

> **property**:
> 1) Something owned; a possession. ...
> 4) a. A characteristic trait or peculiarity, especially one serving to define or describe its possessor. b. A characteristic attribute possessed by all members of a class. See synonyms at **quality**.
> 5) A special capability or power; a virtue.

If you think of a quantum state as a property owned by the system of which it is about, then you—Charlie Bennett in particular—are obliged to continue propagating this thing you told us was a misconception. At the completion of Alice's measurement, there is a new quantum state for Bob's system. If the quantum state is interpreted as a "special capability or power" for the system at Bob's end, then you cannot get around the conception that Alice's twiddle caused a change to something localized way over there.

And that, it seems to me is dangerous business. What I am saying has nothing to do with hidden variables. It just has to do with the word "property". The trouble only has to do with the idea of a quantum state as a kind of feature possessed by a quantum system.

If you want to think of the quantum state as a property of something, it seems to me the best you can do is speak of it as a property of Alice's head (or Bob's head, or whoever's). For, the quantum state represents the *predictions* she can make about measurements upon the system in question. Similarly I could say all the same things about entanglement.

To my view, toying with the idea that a quantum state is a property, is to toy with a kind of pantheism or anthropomorphism that my materialist mind won't tolerate. Do all rocks have souls? You'd laugh at that, but it seems to me that's about what you're attempting to do in thinking of the quantum state as something possessed by the system itself. "That rock judges his chances of reacting to the measurement device to be such and such!" What could be more anthropocentric than that?

Bennettism 6: *"Properties" can reasonably be taken to be much broader than hidden variables, and may include all sorts of conditional and post-selected behaviors, eg "how a system would behave if I measured its Z spin component, after having watched my favorite horse lose at the race track."*

But, as I see it, conditional properties fare no better than the nonrelational type in this regard. At the completion of Alice's measurement, there is a new quantum state for Bob's system, and thinking of it as a conditional property in Charlie's sense still means something physical changed at Bob's end. This is because the system did not have that particular (conditional) property before Alice's action.

The main point is this: Whatever a property for a quantum system is, it should not be something changeable by someone's twiddles far, far away. For instance, take the dimensionality posited for a quantum system by the ascription of a Hilbert space to it. I think this is a perfectly good candidate for a property of a quantum system; it's one I would endorse. For, once set, there is nothing Alice can do at a far away location to change it.

So there, quantum systems do have properties—or at least I'm willing to bet they do. It is just that the properties do not include among them "the" quantum state ... and it is the quantum state that is transferred in the process of teleportation.

What I find miraculous about teleportation is that Victor (the guy who ascribes the original, unknown state) can transfer his predictions from one physical system to another at the cost of only two bits of physical action on the target system. In that sense, it is Victor's description that is teleported from one system to the other with almost nothing whatsoever traveling in between. But that'll never make it into a dictionary.

Of course, this is an ongoing debate between Charlie and me. I'll paste below one piece of our correspondence that has to do with the present conversation—it itself focused on the issue of properties. [See note to Bennett dated, 25 April 2002.]

By the way, I agree with Bill that the word "object" is better than "system." *American Heritage* writes this in its first two definitions for the word:

> **object**:
> 1) A material thing.
> 2) A focus of attention, feeling, thought, or action.

14 May 2002, to Smolin-Bennett and others, "Chris's World"

Good to see you in the morass.

Chris's world: It's a funny place with all these fancy words like "ascription" and "posit" to remind us that there's a head in the background of every quantum state, but there's no instantaneous

action at a distance there—no one would have ever thought there might be. I know the language drives Charlie bonkers, and it probably drives you bonkers too. But mostly the complaints just remind me of what I used to hear in my hometown in Texas when the seatbelt law was first enacted. "Why that Majatek boy was thrown clean from the car! Not a scratch on him. If he'd have had a seatbelt on, his head would have been nothin' but mush now; we'd be at his funeral today. Damned politicians puttin' their noses into places where they ain't got no business."

So you see I view the language as a safety measure—one that I think will allow us to drive farther, longer, and ultimately faster. But first you've got to learn how to drive with a belt on.

Smolinism 1 : *One is naturally forced into Chris's world of saying all that changes is what people predict about the state, but that's a property properly defined as above.*

Chris would never say this. When I write down a quantum state, I think of it as my judgment or prediction for which of one or another measurement outcome will turn up. I don't predict things "about the state." I don't know what the TRUE state is or could be, and thus my ignorance cannot be about it. In contrast, I would say the ignorance is always about further measurement outcomes. Or a pleasing picture is that one can ground the ignorance with respect to a single device sitting in the National Bureau of Standards if one wishes. See Sections 4.2 and 6.1 of the fat paper I put on `quant-ph` last week (`quant-ph/0205039`).

That said, I also see the suppression of another crucial issue here—one that also flies in the face of the word property, especially with your emphasis on DEFINE in definition 4. If you will allow me to call a mixed density operator a quantum state, then I know even you will agree that there is no unique quantum state for a system. Thus it cannot be a property. So, I think you'll be left with being only willing to call a pure state a property. So be it. (I wouldn't do that, but I'll let you do that for the time being.) But now, let's go back to Alice and Bob. By your account, first Bob's system has no property, then Alice measures her system, and—Zing!—now Bob's system does have a property. I.e., first it has no quantum state, then it does. How does the system know it ought to have that property if there's no action at a distance? Alternatively, if it doesn't know it, why call it a property of the system?

It is because of that conundrum that it seems to me to be more fruitful to just give up thinking of the quantum state as a property of the system it targets.

I think a good analogy can be found in *classical* information theory. A homework problem in a textbook gives you a discrete memoryless channel by specifying the transition probabilities $p(y|x)$. Then it asks that you calculate the channel's capacity. One goes through all the work and gets a number. From that, one starts to get the feeling that the capacity can be an objective property for a real physical channel. Why else would it take so much work to calculate something if it weren't real? But it can't be real in any absolute sense. For, with respect to a Laplacean demon, there is never any noise in the channel at all; he can predict which bits will be flipped and which won't. The point is, the capacity is only objective WITH RESPECT TO a *subjective* judgment $p(y|x)$. Similarly, I would say with all quantum states. Just because a textbook says calculate such and such a property—the entanglement of formation, the distillable entanglement, or whatnot—of a quantum state, one finds that one gets into the same habit of thinking those properties have no subjective component.

You can do so, but then it's your burden to explain to young school children and journalists—who think of a property as something like the color red, a ball either has it or doesn't—what changes in Bob's system when Alice performs a measurement.

15 May 2002, to Shimony, "Bayes, POVMs, Reality"

It was good talking to you this weekend. If you wouldn't mind committing your story of meeting de Finetti to email, I will see to it that it is archived forever in one of my samizdats. (See my webpage; link below.)

Also, if you could send me your mailing address, I will send you copies of my two new papers on Bayesianism, POVMs, and good candidates for quantum reality. Alternatively, if you are accustomed to downloading things from the quant-ph archive, here are the links:

1. "Quantum Mechanics as Quantum Information (and only a little more),"
   `http://xxx.lanl.gov/abs/quant-ph/0205039`

2. "The Anti-Växjö Interpretation of Quantum Mechanics,"
   `http://xxx.lanl.gov/abs/quant-ph/0204146`

I wrote them both in an attempt to be entertaining. I hope you find them so.

Abner's Reply

> Here is what I recall of my conversation with de Finetti. In 1971 the 3rd International Congress of Logic, Methodology, and Philosophy of Science was held in Bucharest (where I had the dubious pleasure of meeting Caesescu in the receiving line at the Palace of Ministers; what an ugly hard face he had.) De Finetti was there, and I believe that I introduced myself, saying that I had some questions. The only question I recall was why he didn't use the strong version of coherence. He said that he was aware of the option of using it rather than the weak version, but he didn't like the consequence of the substitution: namely that $C(h/e) = 1$ only if e entails h. This consequence is part of my (2') in Sect. 5 of my "Coherence and the Axioms of Confirmation", p. 136 in vol. I of my *Search for a Naturalistic World View*. I don't recall the details of his objection to this principle, but I think he said that it would cause trouble if one had a nondenumerable set of mutually exclusive possible hypothesis, as in the case of probability on continua. I vaguely recall agreeing with him that there would be a problem, because I discuss the problem on pp. 137–140, op. cit. I vaguely recall saying (or maybe just thinking) that epistemic probability doesn't apply well to nondenumerable sets of hypotheses, but thinking that the propensity interpretation of probability, usable in stat. mechanics, could properly deal with nondenumerable sets of possible outcomes. De Finetti surely would not have liked this discrimination, since he believed that the only clear sense of probability was epistemic, and in particular personalist. In my later paper, "Scientific Inference", op. cit. I suggest pragmatically reasonable strategies for dividing the entire set of hypotheses into as denumerable set, or even a finite set, by properly lumping subsets of hypotheses. This strategy seems to me in the right direction, partly because it is part of a program taking the Bayesian formalism as only a framework, which has to be supplemented by pragmatic and by a posteriori considerations.

16 May 2002, to Smolin-Bennett and others, "King Broccoli"

I'll close with this statement. But then after your rebuttal—if you care to make one—we should probably take this offline. I suspect no one cares to explore the issue further (not even you and not even me). But you said something so nicely, I thought I should emphasize it.

Smolinism 2:

> Chris said: [F]irst Bob's system has no property, then Alice measures her system, and—Zing!—now Bob's system does have a property. I.e., first it has no quantum state, then it does. How does the system know it ought to have that property if there's no action at a distance? Alternatively, if it doesn't know it, why call it a property of the system?

It always had the property that IF Alice measured one thing, then it would behave as state ϕ. Conditional properties like that do, of course, imply a sort of action at a distance, but what's so bad about that? It's not the sort of violate-the-speed-of-light action at a distance that we should be concerned about. In the end, it means just what you want it to, except for the word property. To use a classical example, suppose I have a box that comes from the factory with either a red ball or a blue ball in it. Surely one can say that a property of the box is that it has either a red ball or a blue ball in it—that's the entire definition of the box. Now if I call up the factory and they tell me what color the ball actually is, never touching the box, does the box change? Was I wrong to call that other thing a property? But everyone understands what's going on. Entanglement doesn't really bring in anything new here. (Alternatively you could formulate it like Charlie sometimes does and say well, it was ALWAYS a red ball, but the measurement result travels backwards though time and fixes things up so it is ok to say "always," but I'm sure you won't care for that).

You said that perfectly, and, of course, I especially liked the concession that this choice of words does entail a kind of uninteresting action-at-a-distance ... but action-at-a-distance nevertheless. Maybe my point is just, how is the non-quantum-practitioner supposed to know where to draw the line between the uninteresting and the interesting versions of the effect? Between the science and the science-fiction? How will he ever be able to shake the nagging feeling that we ought to be able to harness the uninteresting version and turn it into the interesting one? (For that matter, how will the quantum practitioner? Nicolas Gisin comes to mind.)

About the particulars of your box example, I think the common man would be hard pressed to take the textbook definition of a problem—like the one you describe above—as a property possessed by a physical system. How can the box containing the red ball know I've embedded it in a problem where the possible colors are {red, blue}, or instead in a problem where they are {red, blue, green}? People think of physical properties as the things physical systems carry around with them independently of the rest of the world. You may say that this is a limiting conception of the word "property," but I'm pretty sure it's the conception most people use. They would call the set {red, blue} a property of the problem you've defined, not of the system. Within classical physics, they would say the ball has whatever properties it really has (say, red OR blue) ... and it is the physicist's task to figure out which of the two it is.

The standard retort you and Charlie give me is that I am doing nothing more than encumbering the language by saying "a quantum state is ascribed to a system" rather than saying "a quantum state is possessed by a system." But, come on, the word count is the same. It is not that I am encumbering the language; it is just that I am beating on a prejudice you don't want to let go of. Or maybe to be more conciliatory, it is that you cannot imagine that this kind of language could ever be useful, whereas you think there are loads of examples where your own language has led to triumph. But I think Charlie's talk the other day about the public's perception of quantum teleportation as a kind of instantaneous action-at-a-distance (in the science-fiction sense) is a case in point.

Things are only interesting or noninteresting with respect to a context. I think there is a sense in which quantum teleportation is less interesting with respect to the conception that a quantum state captures a state of knowledge rather than a property possessed. What could be less interesting

to say than that, "Quantum teleportation is the transference of one's predictions about one object onto another object that has never interacted with the first"? Maybe it's only this that's keeping you in the old bounds. Is it that if you keep a little bit of the science-fiction imagery alive, it'll help fuel the physics?

"In any case, none of this matters for doing physics," you say. But I think it does in the long run. (Certainly, you've got to concede that there's something that fuels me—and it hurts to think that you might think it is nothing more than irrationality.) When Charlie sent me the picture of a skunk cabbage in reply to these very issues, I found myself thinking of King Broccoli. The story goes that one day, by divine providence, it came to King Broccoli that broccoli, the vegetable, his namesake, actually tastes good. Good in a way that hitherto only gods and angels had known. Every child who had ever said, "Yuck, I don't like broccoli; it tastes awful!" was simply wrong ... or at least that's what the king realized. King Broccoli, being the head of state, decided to do something about it. Henceforward, all gardens in the kingdom should have a patch devoted solely to broccoli. It really wasn't much of a burden on the national product (except, perhaps, for the psychiatrists who had to treat all the movie stars who had never felt fulfilled in their broccoli experience). But think of the diversity of vegetables the kingdom might have raised if its citizens hadn't been encumbered with the king's notion that broccoli had an objective, but never verifiable, taste?

Moral? Maybe there's none. But it is a documented fact that the Kingdom of Broccoli eventually fell and was replaced by a liberal democracy (where the ideals rather than the particulars have an objective status).

Everyone in this mailing list knows by now—though Charlie and John seem to keep forgetting this, or maybe they've never let it sink in—that I think quantum mechanics is just the hint of something much deeper, some fantastic physics yet to come. But I also don't think we will ever stumble upon that physics until we truly get rid of our classical prejudices: seeing quantum states as "properties"—it seems to me—is one of these.

Signing off,

Chris

My own disclaimer: Though I implied above that quantum teleportation becomes less interesting within a subjective conception of the quantum state, I think effects like quantum cryptography become *more* interesting from this view. So there is a tradeoff.

17 May 2002, to Rudolph, "No Nasty"

I won't be nasty; I enjoyed your note.

Your thinking has a lot of the flavor of the paper: A. Peres and W. H. Zurek, "Is Quantum Theory Universally Valid?," Am. J. Phys. **50**, 807 (1982). I'm not sure if you're aware of it. Asher once told me that he himself still likes the paper a lot, but Zurek basically disavows it now.

Rudolphism 1: *what I'm trying to understand is the physics analogue of Turing's construction – what is it that I, a regular physicist, am doing in my interactions with the world and my construction of theories to explain those interactions?*

I like this question a lot, actually. Below is some of my own attempt to ask the same thing. In this regard, I suggest you read Richard Rorty's book, *Philosophy and Social Hope*. I think if I were to read it again, I would realize how much of an influence it's had on my own thoughts. In fact, my Anti-Växjö paper may be little more than a condensation of it.

I guess if you were to ask me now—in my present state of mind—I would say we are doing more than constructing theories. We are constructing the world (in part). See other note below. But beyond that, I don't know how to say more presently.

Keep thinking about your question! Don't listen to Spekkens.

By the way, I won't be around for essentially a month: Today, I leave for Texas until Thursday. Then the Saturday after that I'm off to Brisbane for a 3 weeks. I return to Bell Labs June 17. (I'm still looking for volunteers on the fence.)

17 May 2002, to Wootters, "Dueling Banjos"

Thanks for the Hartshorne quote. I had forgotten that I had sort of agreed with that a bit. Here's the way Martin Gardner put it in his essay "Why I Am Not a Solipsist":

> In this book I use the term "realism" in the broad sense of a belief in the reality of something (the nature of which we leave in limbo) that is behind the phaneron, and which generates the phaneron and its weird regularities. This something is independent of human minds in the sense that it existed before there were human minds, and would exist if the human race vanished. I am not here concerned with realism as a view opposed to idealism, or realism in the Platonic sense of a view opposed to nominalism or conceptualism. As I shall use the word it is clear that even Berkeley and Royce were realists. The term of contrast is not "idealism" but "subjectivism."

(The phaneron, by the way, was Peirce's term for "the world of our experience—the totality of all we see, hear, taste, touch, feel, and smell.")

Thus, in making the transition from the first paragraph to the first sentence of the second paragraph in the excerpt from my old letter below, I was making a non sequitur.

Let me ask you this about your "idealistic current." Does it run counter to what my first paragraph below says and what Gardner says above? I guess that's the main point for me.

I suppose if I were to start to label things, then this thing I was telling you about the other day—"the sexual interpretation of quantum mechanics (SIQM)"—would be a kind of dualistic theory. I said it metaphorically this way: When things bang together, something is created that is greater than the sum of the parts. Or again: When things—that's the materialistic aspect—bang together, something is created—that's the mentalistic aspect, for it is like an act of the will or a decision. But that's just a thought that's hitting me at 4:00 in the morning. (So trust it less than even the usual things that come out of my mouth.) I hadn't thought about it in this way before, and I'm not sure I want to continue to thinking about it this way. In general, I don't like dualisms. (Though even saying that is paradoxical; for I think I like "pluralisms" in the sense of James.) The excerpt from an old letter far below gives a slightly longer introduction to the idea of the SIQM.

17 May 2002, to Mermin-Bennett, "Morning Coffee"

Boy you got me into a stink, didn't you, by getting me into that word debate! If you hadn't done that, I might have had some time to answer your other questions this week.

Below for your continued amusement, I'll include a side conversation I had with Charlie. Maybe this debate is at least edifying in some ways. Charlie sees me as hair-splitting; I see Charlie as being half-hearted and inconsistent.

I just read these words in a Martin Gardner article:

A third aspect of aesthetic theory that bores me even more are all those tiresome disputes, in book after book, about whether aesthetic values are subjective or objective. Here the situation is not quite the same as that of truth. In previous chapters I have argued that the least confusing way to talk about truth is to assume that the world and its structure are not mind-dependent. But beauty, so far as humanity is concerned (we will not consider what beauty may mean to birds or apes, to creatures on other planets, or to gods), obviously requires a human mind. Where is the red of an apple? As I have said, it is in the mind if by red we mean the sensation of red. It is on the apple if by red we mean the structure of a surface that reflects a pattern of visible light which causes a mental sensation of red.

I see no difference between this antique quibble and the question of whether beauty (however defined) is a property of an art object or a sensation in a brain. If by beauty you mean the pleasure aroused by a beautiful object, of course it is subjective. If by beauty you mean the structure of an object capable of arousing aesthetic pleasure, then the beauty is a part of the object. Or you may prefer a third approach and ground beauty in the combined dynamic structure produced by the interaction of an object and a mind. It is all such a weary waste of words. The last approach is the one taken by John Dewey in his influential book Art as Experience. Although I found fault with Deweys attempt to redefine truth in pragmatic terms, I find his approach to aesthetics (essentially the same, by the way, as Aristotle's) a sensible way of speaking. Again, it is not a question of Dewey being right or wrong. It is a question of the most useful way to talk about aesthetic values.

And I find that I agree with most every word of this (except for the part about the Jamesian-Deweyan theory of truth). But the issue at stake with quantum mechanics goes much deeper than this, and it annoys me to no end that our friend Charlie lumps me in with the art critics.

If you have any words of wisdom that could take a little fire out of our relations, I'd love to hear them. Maybe I'll CC this note to Charlie too.

By the way, this was not intended to be the subject of this note. I wrote instead to tell you that I have now posted the slides from my talk at your party on my website. Maybe you'll enjoy seeing the second half, if not hearing it.

Charlie,

Bennettism 7: *Surely you must have something more sensible to say than some of us who have spoken.*

It would have been more neutral to say, "than those of us who have spoken." If I were trying to read between the lines, I might be tempted to write a note just like this one. But I'll refrain from reading between the lines.

Chris

Dear Chris,

Sorry. Point well taken. You are a perpetual stimulus to me, if not always in the ways you hope, and I would miss it terribly if you stopped. I should be more grateful. Like the other day, when John was discussing your automotive metaphor, you inspired me to think that all the cautions you would have us take against quantum misconceptions are—for me, if I did them—like driving a car with the parking brake on all the time.

-CHB

Charlie,

Bennettism 8: *Like the other day, when John was discussing your automotive metaphor, you inspired me to think that all the cautions you would have us take against quantum misconceptions are—for me, if I did them—like driving a car with the parking brake on all the time.*

Then it seems to me you should at least be consistent in your behavior. I interpreted your talk Sunday as a genuine concern for the perceptions the masses hold about some of our favorite quantum effects. Do you have a concern, or do you not? Is it that the world really has some kind of instantaneous action at a distance—like John's note yesterday supported—and we're just not allowed to say the phrase in polite company, or does it not? If the world does not, then so be it. But if it does, why should we try to so hard to delete the phrase from polite conversation? The main point I always wonder is how well you really have these issues worked out in your own head. If you care about misconceptions, then care about them—I say in this slightly grumbling state—and if you don't, then don't. But I have trouble understanding your mix of halfheartisms.

Chris

Dear Chris,

I think the world does not have instantaneous action at a distance, and it's important to find ways of speaking that do not encourage the frequent misconception that entanglement provides a means of faster than light communication. I am less interested, and think you should be less interested, in what seem to me to be hair-splitting arguments about "properties" or "changes in one's state of knowledge" that merely reflect different ways of talking about situations in which we entirely agree about the predictions of outcomes for any experiment. That seems a matter of aesthetics only, and the protections you find so reassuring I find merely annoying, like a parking brake, since they complicate my language and don't prevent me from making any wrong predictions, just what you would call wrong or fuzzy interpretations. I would say that if a way of thinking such as my own does not lead to wrong predictions, then it is not fuzzy in any serious way. Or to put in another way, the distinctions you would have me make because they seem so important to you are in fact less real than the wave function you don't want me to believe in.

Talking with you first thing in the morning is way better than coffee.

-CHB

17 May 2002, to Schack, "More Balking"

Schackcosm 68:

Chris said:

> Rüdiger said, "I am still convinced, despite your severe scolding, that a priori, without looking at the consequences, strong and weak coherence are similarly compelling axioms."
>
> I don't see how you can make this distinction. It seems to me, axioms are only compelling or not compelling insofar as their consequences.

Then why would there be a difference between postulating the Kolmogorov axioms and postulating coherence? Why would we bother with the Dutch book arguments? Why do you bother looking for an information-theoretical reason for the quantum axioms?

It seems to me that the compelling reason to accept the axiom of (ordinary) coherence is that "whatever the consequences are, I certainly do not want to violate coherence, because I don't want to hand over money". It seems to me that a discussion of whether strong coherence is a compelling axiom has to be carried out at this level.

But that is not what is at issue. Strong coherence and regular coherence are *two different theories*, and a theory's worth can only be assessed by looking at the whole thing.

It seems to me a very simple issue. In the theory of strong coherence, there is a kind of normative behavior for all agents that regular coherence does not have. The theory says, in effect, "Thou shalt not tell a lie." If an agent writes down $p = 1$, then in the rigid world of strong coherence, I can trust his statement to be a reflection of his true inner belief. The world of regular coherence does not have that.

You can say that one theory is just more of the same with respect to the old theory. But that is to only look at one aspect of the problem. In another aspect of it, the two theories become qualitatively different.

There is a sense in which my mother-in-law is just a faster version of Kiki at times. That is, she's just more of the same when it comes to cooking, artistry, and a couple of other aspects. But I fell in love with Kiki, not my mother-in-law. This is because when I move past a few isolated features of my mother-in-law, I discover she is qualitatively different from the woman I love.

How can you reject that as a reasonable line of thought?

If you thought at the outset that probabilities could be proper properties of things—like Shimony does—then you might indeed accept strong coherence. For strong coherence grounds the very meaning of $p = 1$. A "$p = 1$"–statement is a TRUE statement. But if you start out from a strongly subjectivist stance on probability, then the only leg you have to stand on is the kind of argument you and Carl give: It is just more of the same of regular coherence; it is just slightly more sensible or cautious betting behavior. But I counter that by saying, carried through consistently and without exception, strong coherence makes it nigh on impossible to embed one's particular bets about particular events in the larger framework of all bets and all events.

It is an issue worthy of debate, I agree, whether my reason is a decent reason to reject strong coherence. (I think it is, clearly.) But it's not fair to say in distinguishing horses and zebras, we should all close our eyes and focus on tactile differences, eschewing all the visual information around us.

17 May 2002, to Wootters, "The Divinity that Breathes Life"

Wootterism 5: *I know you will say that your realism is not the same as that of classical physics. But in what way, ultimately, is it not the same, other than by the absence of determinism? I think indeterminism is crucial, but I think one needs the further step of restoring the subject-object relation as fundamental.*

To the extent that I could write something like the passages below, I don't think so. [See note to Renes, dated 20 November 2001.] Or, say, to the extent that I find myself liking this little quote, "I forbade any simulacrum in the temples because the divinity that breathes life into nature cannot be represented," I don't think so. I think there's more than simple indeterminism in my forming view; there is something lower level than determinism and indeterminism both.

But what I am not settled on is whether the "divinity that breathes life" is the subject-object relation. I will think hard about your point, and I will think harder about the impression my writings give.

20 May 2002, to Smolin, "Dizzy in Texas"

Smolinism 3 : *Now you guys are trying too hard not to think like Chris. Nothing is magic about the "instantaneous" first stage of teleportation, only knowledge is being changed. If I have two boxes, one with a ball in it and one without and send one far away, when I open the remaining one I instantly know if the ball is in the other one. Big deal. When I measure something, I find something out. Wow!*

I shouldn't let your note encourage me. (The sober side of myself won't.) But there is a way to make the sum content of all quantum measurements look exactly like your analogy above, even in a formal way. The mathematics may interest you, even if not the philosophy. You can find it in Sections 4.2 and 6.1 of the paper I put on `quant-ph` last week (`quant-ph/0205039`). Or you can see it sketched quickly in the Merminfest talk posted at my website.

The trick is to represent the quantum state as a probability distribution with respect to a fixed, fiducial informationally complete POVM. That is, one imagines a "standard quantum measurement device" sitting in the National Bureau of Standards beside the standard meter and the standard kilogram, and the quantum state captures nothing more than one's judgment for how the device will react if one were to throw one's quantum system into it. With that picture in mind, all a regular, everyday measurement does is update one's judgment concerning the outcomes of the standard measurement. The twist that comes with quantum mechanics (over regular probability theory) is that the update rule is Bayes' rule plus a little more ... it's not just Bayes' rule full stop.

25 May 2002, to Nielsen, "Fun with Feyerabend"

I'm just putting some notes into my computer before I set off for travel again. At your recommendation, I've paid a little more attention to Paul Feyerabend. The quotations below summarize a little of what I see as likable and usable in his philosophy.

See you tomorrow.

From: P. Feyerabend, *Conquest of Abundance: A Tale of Abstraction versus the Richness of Being*, edited by B. Terpstra (U. Chicago Press, Chicago, 1999).

> *Humans as Sculptors of Reality.* According to the first assumption, our ways of thinking and speaking are products of idiosyncratic historical developments. Common sense and science both conceal this situation. For example, they say (second assumption) that atoms existed long before they were found. This explains why the projection received a response, but overlooks that vastly different projections did not remain unanswered.

A better way of telling the story is the following. Scientists, being equipped with a complex organism and embedded in constantly changing physical and social surroundings, used ideas and actions (and, much later, equipment up to and including industrial complexes such as CERN) to *manufacture*, first, metaphysical atoms, then, crude physical atoms, and, finally, complex systems of elementary particles out of a material that did not contain these elements but could be shaped into them. Scientists, according to this account, are sculptors of reality—but sculptors in a special sense. They not merely *act causally* upon the world (though they do that, too, and they have to if they want to "discover" new entities); they also *create semantic conditions* engendering strong inferences from known effects to novel projections and, conversely, from the projections to testable effects. We have here the same dichotomy of descriptions which Bohr introduced in his analysis of the case of Einstein, Podolsky and Rosen. Every individual, group, and culture tries to arrive at an equilibrium between the entities it posits and leading beliefs, needs, expectations, and ways of arguing. The separability assumption arises in special cases (traditions, cultures); it is not a condition (to be) satisfied by all, and it certainly is not a sound basis for epistemology. Altogether, the dichotomy subjective/objective and the corresponding dichotomy between descriptions and constructions are much too naive to guide our ideas about the nature and the implications of knowledge claims.

I do not assert that any combined causal-semantic action will lead to a well-articulated and livable world. The material humans (and, for that matter, also dogs and monkeys) face must be approached in the right way. It *offers resistance*; some constructions (some incipient cultures—cargo cults, for example) find no point of attack in it and simply collapse. On the other hand, *this material is more pliable than is commonly assumed*. Molding it in one way (history of technology leading up to a technologically streamlined environment and large research cities such as CERN), we get elementary particles; proceeding in another, we get a nature that is alive and full of Gods. Even the "discovery" of America, which I used to support the separability assumption, allowed some leeway, as is shown by Edmondo O'Gorman's fascinating study, *The Invention of America*. Science certainly is not the only source of reliable ontological information.

It is important to read these statements in the right way. They are not the sketch of a new theory of knowledge which explains the relation between humans and the world and provides a philosophical grounding for whatever discoveries are being made. Taking the historical character of knowledge seriously means rejecting any such attempt. We can describe the results we have obtained (though the description will always be fatally incomplete), we can comment on some similarities and differences that have come to our attention, and we can even try to explain what we found in the course of a particular approach "from the inside," i.e., using the practical and conceptual means provided by the approach (the theory of evolution, evolutionary epistemology, and modern cosmology belong in this category). We can tell many interesting *stories*. We cannot explain, however, how the chosen approach is related to the world and why it is successful, in terms of the world. This would mean knowing the results of all possible approaches or, what amounts to the same, we would know the history of the world before the world has come to an end.

And yet we cannot do without scientific know-how. Our world has been transformed by the material, spiritual, and intellectual impact of science and science-based technologies. Its reaction to the transformation (and a strange reaction it is!) is that we are stuck in a scientific environment. We need scientists, engineers, scientifically inclined

philosophers, sociologists, etc., to deal with the consequences. My point is that these consequences are not grounded in an "objective" nature, but come from a complicated interplay between an unknown and relatively pliable material and researchers who affect and are affected and changed by the material which, after all, is the material from which they have been shaped. It is not therefore easier to remove the results. The "subjective" side of knowledge, being inextricably intertwined with its material manifestations, cannot be just blown away. Far from merely stating what is already there, it created conditions of existence, a world corresponding to these conditions and a life that is adapted to this world; all three now support or "establish" the conjectures that led to them. Still, a look at history shows that this world is not a static world populated by thinking (and publishing) ants who, crawling all over its crevices, gradually discover its features without affecting them in anyway. It is a dynamical and multifaceted Being which influences and reflects the activity of its explorers. It was once full of Gods; it then became a drab material world; and it can be changed again, if its inhabitants have the determination, the intelligence, and the heart to take the necessary steps.

and

> Pauli's views have much in common with the general picture that emerged from Aristotle's principle. In this picture we start with a world (which I shall call the primal world, or Being) which behaves in its own way and not necessarily in accordance with any one of the laws that have been discovered by scientists. (Here we still have an element of realism.) Humans are part of the primal world, not detached aliens, and they are subjected to its whims: Being can send scientists on a wild-goose chase—for centuries. On the other hand, it permits partial independence ... and it provides some of those acting independently (not all of them!) with *manifest worlds* they can expand, explore, and survive in (manifest worlds are in many respects like ecological niches). Inhabitants of a particular manifest world often identify it with Being. They thereby turn local problems into cosmic disasters. But the manifest worlds themselves demonstrate their fragmentary character; they harbor events which should not be there and which are classified away with some embarrassment (example: the separation of the arts and the sciences). The transition from one manifest world to another cannot be described in either except by excising large regions originally thought to be real—a good case for applying the notion of complementarity. Bell's request that a fundamental theory should not contain any reference to observation is satisfied, but trivially so. Being as it is, independently of any kind of approach, can never be known, which means that really fundamental theories don't exist.

25 May 2002, to Comer, "Becoming William James"

A thought struck me today as I was flying from Newark to LA reading Lloyd Morris's book *William James: The Message of a Modern Mind*. It is this: I think there's nothing I presently want more than to become a modern (quantum) incarnation of William James.

Let me leave you with a lovely quote from the end of Chapter 2.

> He conceived the individual's life, and all social progress, as a form of perpetual experiment. But he did not preach reckless faith. The "will to believe" for which he argued is best defined as courage weighted with responsibility. Contingency signifies

that no precaution can absolutely eliminate all hazard of shipwreck. The individual must take everything into account that may tell against his success, and make every possible provision to minimize disaster in the event of his failure. But having done so, he must act. And in this circumstance, James preached the right of the individual to indulge his personal faith at his personal risk. The part of wisdom would always be to believe what is in the line of one's needs, for only by such belief is the need fulfilled. Over a wide area of existence, possibilities and not finished facts are the realities with which we have actively to deal. So James argued, and pointed out that "as the essence of courage is to stake one's life on a possibility, so the essence of faith is to believe that the possibility exists." But his doctrine subordinated faith to action, for the real utility of faith is to make action genuinely dynamic. "These, then, are my last words to you," he told a group of Harvard students. "Be not afraid of life. Believe that life *is* worth living, and your belief will help create the fact."

Now I'm off to Brisbane. For my own self, I believe it is worth becoming William James.

27 May 2002, to Svozil, "Late, but Never Too Late?"

I hope you will excuse my horribly late reply to your letter of 4/19/02. I have just gone through a month of hell of traveling almost constantly, and I have gotten pathetically behind in all my correspondence. (Just check with Johann Summhammer there; he is also in my queue! I'm hoping to get to him later in the day, if not tomorrow.)

Thanks for the continued interest in having me around for the quantum structures meeting. Regrettably, I think I'm going to have to bow out of the possibility of coming. Presently I'm in Australia, separated from my wife and children for a month, and this trip is making me realize that I shouldn't take on any further travel than the stuff I'm already committed to for the summer. It's a shame really, because I am getting ever more involved in IQSA kinds of ideas, and it would be a great opportunity for me to learn a lot more about what is already "out there" mathematically ... just waiting for me to plaster over with some words of Copenhagenish flavor. Beyond that though, I would love to have a chance to express some of my point of view to that audience. I think my talk at the quantum structures session at the AMS meeting in Atlanta two months ago or so went exceedingly well in that regard, and I found myself really enjoying conversations with Dave Foulis, for instance, afterwards.

I finally had a chance to read your paper "What Could Be More Practical than a Good Interpretation?," by the way. There are certainly large parts of it I agree with—I don't know if that will shock you or not. However, there was more to my paper with Peres than just the title! I do get a little shocked when I find people reading the title of the paper as its sum content. Here's the way I put it to Philippe Grangier a few months ago:

> **Grangierisme 6:** *By the way I also disagree with your point of view that "Quantum Theory Needs No 'Interpretation'," Phys. Today* **53***(3), 70 (2000). The fact that a physical theory ALWAYS needs an interpretation is in my opinion a central difference between physics and mathematics.*
>
> You won't find a disagreement with me here. The title and closing sentence of that paper were meant to be tongue-in-cheek plays on something Rudolf Peierls once said: "The Copenhagen interpretation *is* quantum mechanics." The whole paper is very

definitely about an interpretation, and why one does not need to go any further than it to make sense of quantum mechanics as it stands. My paper `quant-ph/0106166` and the large (more personal) collection `quant-ph/105039` is about going the next step, i.e., what to do once we have established the belief that quantum states are states of knowledge.

When we do finally dig up an ontology underneath quantum mechanics, I'm quite sure it will be an interesting one!

And here was the way I put it to Paul Benioff a year before that:

Benioffer 1: *To me that is an interpretation of QM. Interpretations are what give otherwise empty theories their meaning.*

You're quite right about that. What Asher and I wrote about is indeed a kind of interpretation of the quantum mechanical formalism. The title and the ending words of the article were more for attention-getting than anything else. Also, though, the words were meant to be a small slap in the face to some of the extremes people have gone to (like Everett worlds, Bohm trajectories, and Ghirardi-Rimini-Weber stochastic collapses) just to hold on to a philosophic view that came around long before quantum mechanics was ever heard of. (Talk about people being set in their ways!)

I think the best renditions of my present views can be found in my new papers, `quant-ph/0205039` and `quant-ph/0204146`. I certainly would appreciate any comments you have on those. Especially I would love for you to articulate the weak points you see in the ideas. It is an evolving point of view. And what I want out of the effort is just what you suggest in your paper. I want new calculations, new effects to go search for, new mathematics, and *really* a new world view in total.

Will by chance you be at either the Oviedo, Spain meeting or the QCMC meeting at MIT this summer? Maybe we could talk then?

28 May 2002, to Summhammer, "Anti-Anti-Växjination"

Thank you for the beautiful long note of 4/29/02. Please allow to apologize for taking so long to reply. I've been almost completely incapacitated in my email efforts for a few weeks now by travels, family vacations, company duties, etc. Certainly I appreciate the efforts you take to read my things, and definitely your questions and comments are the best products of that!

Let me comment on a couple of your points.

Summhammerism 5: *But this 'core of the theory' is always tentative. New information, a wider frame of thought, can change it. And yet it is hard to deny that it captures something about that which has been observed. It is like clouds in the sky. For some time they do look like an animal, or the face of a witch, and anyone with eyes will agree, and a few minutes later they are gone. The immutable part here, as well as in physical theories, does seem to be the rules of thought. By the term 'core of the theory' you seem to want to say that they contain a timeless truth. But I have difficulties believing that a physical theory could ever achieve the degree of timeless truth as exhibited, for instance, by mathematical theorems, which are particular expressions of the rules of thought.*

Not quite. I would not say that this thing "the core of a theory" contains a timeless truth. Indeed I tried to be careful to squash that idea when I wrote:

> **Woottersism 6:** *But you obviously also want to say that our theories tell us* something *about reality, even if they are not descriptions of reality.*
>
> I hope you can glean from all the above that I do indeed believe our theories tell us something about reality. But that something is much like what the elephant tells us about reality. Its presence tells us something about the accumulated selective pressures that have arisen up to the present date. A theory to some extent is a statement of history. It is also a statement of our limitations with respect to all the pressures yet seen, or—more carefully—a statement of our limitations with respect to our imaginations for classifying all that we've yet seen. (I for instance, cannot jump off the leaning tower of Pisa unprotected and hope to live; you, for instance, cannot get into your car and hope to push on the accelerator until you are traveling beyond the speed of light.) Finally, to the extent that we the theory users are part of nature, the theory also tells us something about nature in that way.

Thus even the core of the theory is as historical and contingent as the elephant. And just as the elephant could disappear from all historical record, so could the theory and, with it, its core.

Summhammerism 6: *I did not understand what you meant with the last paragraph on page 16 ('However, we would never haven gotten to this stage ....') Sounds as if you could envision that we could set the laws according to our wishes. Reminded me of an old view of evolution: The will (the basic entity) wanted a hand for this particular species, and so it came about ... Sometimes I like this idea, because it permits to look for other patterns and correlations in the history of life than 'mindless' Darwinism.*

Nor did I understand it really. But, yes, I suppose I'm imagining that we might have some control in shaping the "laws" of physics. However—and this is important—that control should be no stronger than the kind of control we might have for shaping a species with genetic engineering. The species has to be able to survive on its own after being produced in a trial. If it can't survive on its own, then one would be loath to call the monstrosity so produced a species to begin with. And so too with what I imagine for this malleable universe.

28 May 2002, to Wootters, "Little Toes and a World of Experience"

I'm finally down in Australia, with a little time to think. I'll be down here for three weeks. Kiki and the kids are in Munich, visiting Kiki's mom and dad.

If you don't mind, I'd like to ask you a couple of questions about your last letter.

Woottersism 7: *I think the sort of world envisioned by classical physics is in fact impossible. If we really understood what it takes to make a world exist* really*, and not just on paper, I think we would see that one needs the subject-object relation in order to hold things together. This is not to say that one needs dualism. In Whitehead's system, everything is both subject and object, depending on the point of view. There is no dualistic separation into two kinds of entity. But everything is*

related to something else as subject, and everything is also related to something else as object. In classical physics, there is no such relation. As Schrodinger points out, from the very beginning we eliminate the very notion of an experiencing subject. I agree with Schrodinger that this is too extreme an abstraction. In making this abstraction, we have removed something essential from our view of the world.

I suppose I too have a gut feeling that your first sentence above is on the right track, but I wonder if you have an actual argument in mind for supporting the case? Also let me ask you this. Would you say the same thing about a world governed by quantum mechanics? Or do you think the quantum world differs from the classical world in this respect? Maybe, in a refinement of this, let me ask the same thing not about the quantum world in general, but the quantum world in the vision of the many-worlders, say David Deutsch and Charlie Bennett. The way I understand what they envision for quantum mechanics, it has never struck me as so very different from what I would call classical physics. (It is just now that the world as a whole has this thing called a "state" and it is it that goes along according to some mathematically precise law.)

Woottersism 8: *I would like to think that my view avoids the line of reasoning that takes the mathematical description as the essence of the world. As John Wheeler says, a set of mathematical laws will not "fly" by itself. I think the necessary added ingredient is something like* experience. *And that's what I find in Whitehead's view.*

Would you flesh out this thing you call the "subject-object relation" a little more? What do you mean by it? Let us focus on a simpler system than one usually thinks of as a sentient being. Say a rock or the little toe on my right foot. When we think of these systems in their capacities as subjects, rather than objects, what is it that defines those capacities? What are their characteristic traits?

When I think of a person as a subject, I think of him carrying around sets of probability distributions for this and that. That is, I think of a subject as something that can carry beliefs. However, I guess, I have a hard time thinking of a rock or my toe as carrying around beliefs. I also think of a subject as something that can play an active role in shaping other parts of the world because of those beliefs.

However, maybe you mean something completely different when you're thinking of a rock as a subject.

Beyond that, let me ask about this word "experience." What do you mean by that? I hope to try out one of Whitehead's shorter books soon but I haven't gotten to that yet. However, I did read another small book on James on my flight over and it got me into some territory that I haven't yet seen of his. Namely, his stuff from *A Pluralistic Universe* and *Radical Empiricism*, neither of which I've read yet. The book I read was *William James: The Message of a Modern Mind*, by Lloyd Morris. It's only 91 pages, and I found it an engaging little thing (at least for giving me a quick overview of all of James's views, even if not the arguments he used for getting there). Anyway, Chapter 5 was titled "A World of Pure Experience" and what he described there seemed to have some overlap with what you expressed in your last note. Morris wrote:

> "My thesis," James declared, "is that if we start with the supposition that there is only one primal stuff or material in the world, a stuff of which everything is composed, and if we call that stuff 'pure experience,' then knowing can easily be explained as a particular sort of relation towards one another into which portions of pure experience may enter. The relation itself is a part of pure experience; one of its terms becomes the subject or bearer of the knowledge, the 'knower,' the other becomes the object known."

> This doctrine is essentially monistic. But it is radically unlike the monistic doctrines of either idealism or materialism, which respectively affirm that mind and matter are the ultimate substance of reality. Pure experience is neither mind nor matter, but is the ground of both. In itself it is, as James asserted, neutral.

Morris also said that this had some feed-in to Whitehead's later thought, but he didn't elaborate. So I'm guessing there will be some overlap between this and what you're thinking about. But still I'm having trouble understanding what all this might mean, especially since I'm having trouble envisioning the mental life of my toe.

Now the other day I said the idea of a random outcome in a quantum measurement might be viewed as having a mentalistic aspect. From the outside, it looks like a random occurrence; however, from the inside one might think it looks like a "decision." But right now, I'm wondering what even I really mean by this.

Thinking of you from way down under.

29 May 2002, to Mermin, "I Think She'll Know"

Remember what the dormouse said; feed your head.

You can tell I'm pretty darned behind in my email. I'm in Australia now, finally with a little time to think. Kiki and the kids are in Munich. Caves, Schack, and I are down here for three weeks doing Bayesiany things.

Merminition 47: *I don't see what your teleportation example (pages 11, 12) adds to ordinary EPR. Aren't all the issues exactly the same if Alice "in her laboratory prepares" the single qubit in (1) that she possesses by an appropriate measurement (to be sure, she can't control which outcome she'll get, but that doesn't seem to be central to your point, or is it?) after which she and only she knows what the outcome of the corresponding yes-no measurement on Bob's qubit will be.*

Yeah, I'd agree that it doesn't add a heck of a lot to the old argument. Mostly I wanted to say something about teleportation: Namely that if Einstein had known about it, then he might have used it to the same devious purposes he did with his old argument.

But still, I guess there were a couple of features of this version that I thought made it a bit cleaner than the old Einstein thing. 1) With the supplementation of *only* two bits of physical action on the part of Bob (i.e., one of four possible unitaries), Alice can put Bob's system into any state she wishes. So, in essence (i.e. up to two bits), there's nothing even random and uncontrollable about the process. 2) In the case of teleportation, even examining the measurement device before and after the measurement will tell you nothing about the posterior state for Bob's system that Alice ends up with. Very literally, only Alice ends up doing some updating. If it's she and only she, why not call the state her knowledge?

29 May 2002, to Mermin, "That Damned von Neumann"

Now let me try to answer more adequately the question you asked after my talk in Ithaca. The point is simply this: Suppose I tell you that I've got a device that measures a standard observable H. How do you know that you should accept my claim? Let's say you do this: You simply give me a supply of a gazillion nonidentical states you've prepared (anyway you wish) in your laboratory. I'll perform my measurement on each of them and report the results I found. If you find that

I'm giving you back outcomes with the (conditional) statistics you expect, then you'll have some warrant for believing I'm really performing the measurement I claim.

Now let me ask you this. Suppose you confirm my measurement to your satisfaction by that method. Do you now have warrant to say anything about the post-measurement state for each of the systems you gave me? The answer is "no" of course. The point is, you need to know more about the particulars of the device.

Now, von Neumann said that for an "ideal" device the post-measurement states for the systems will be eigenstates of the observable. But I claim that is an arbitrary notion of ideal, and Kraus's theory of "effects and operations" backs that up in a kind of technical way. The Kraus theory says that the state change can always be thought of as a collapse in the von Neumann sense PLUS a trace-preserving completely positive map.

Now, the standard quip that is made is that the CPM part of this is just extra noise that didn't need to be there. But again, I claim this statement is arbitrary. Here is a simple counter example. Suppose I perform the measurement H on half of an entangled pair. Then (via the entanglement), I can always think of this measurement as *really* a measurement of some sort on the other half. In fact I can think of it as simply the observable that is the transpose of the original one ... only performed on the second system rather than the first. What could be a more minimally disturbing way of measuring H^{T} on the second system than that? But does von Neumann's collapse postulate hold for this kind of measurement? Try it, and you'll find that it doesn't. The only state changes this kind of measurement can produce is of the pure-refinement kind.

Von Neumann brainwashed the generations in a needless way. Even something so simple and "ideal" as a standard photon detector does not follow the postulate. When there's a click, the photons are absorbed from the mode and the field is now in the $|0\rangle$ state.

I hope that helps make up for my lack of lucidity at your celebration.

29 May 2002, to Skyrms, "More Strict Coherence"

> Brian said: I think that the usual response to Shimony is just that strict coherence isn't a plausible requirement in the usual way of doing probability theory with continuous mathematics.

That's too bad. What I'm really looking for is anything that expresses a dissatisfaction with strict coherence with respect to discrete event spaces. It turns out that enforcing strict versus normal coherence can make a pretty drastic difference for some interpretive problems in (even finite dimensional) quantum mechanics. Unfortunately, with respect to the attractiveness of the conclusion, my coauthors and I have opposite opinions. Thus, I am inclined to require only normal coherence for the agents in our game; whereas they are inclined to require strict coherence.

The main point, even in the classical case, is that strict coherence requires that an agent ascribe probability 1 to an event if and only if he believes the event is a certainty. Whereas under the assumption of normal de Finettian coherence, a probability-1 assignment cannot be used to conclude a belief of certainty on the agent's part. He might be assigning probability 1, not because it reflects his true beliefs, but because it is advantageous for other purposes. That is, with respect to certainty, strict coherence compels an agent to never "tell a lie."

I think that goes too far as concerns a foundation for "rational behavior." Whereas normal coherence appears to me to strike a sweet spot.

In any case, what I am looking for is some confirmation of my troubles in the published literature ... to help me build a case for the inevitable battles I foresee with my coauthors.

By the way, you might be interested know that we have been putting a substantial effort into interpreting quantum probabilities as personalist probabilities. Let me recommend four of our papers to you (along with the web links to get them) in case you're interested. [quant-ph/0205039, quant-ph/0204146, quant-ph/0104088, quant-ph/0106133]

As you climb from bottom to top in this list, you'll find us moving closer and closer to a personalist position. Also, there is a lot of supporting information posted at my website, link below.

31 May 2002, to Wiseman, "Poor Young Duvenhage"

Thanks yourself for your long note in reply to my long note! You can see I'm working with a much larger lag time than you in my emails. Anyway I'm in Australia now, and slowly working off the jetlag and becoming more productive as the days go by.

Let me say a few words on your last note.

Wisemanism 4: *Freedom taken, thank you. And I think you will be stronger if you don't try to prop up your position using another paper which I think goes fundamentally against your views in a number of places.*

If you believe that, you misconstrue my purposes for citing Duvenhage. My citation was to give attention to and encourage this young researcher. Of course, I don't agree with a lot of the paper; but that is beside the point for me. There is one thing I certainly agree with and I think he said it particularly cleanly.

You can find where I said this outright to Duvenhage himself by looking at pages 167 and 168 of my web samizdat, "Quantum States: What the Hell Are They?" (It's in a 28 March 2002 letter; I hope you'll read it.) There you will find me expressing some of the same misgivings as you, though with less detail. As I told him, I would rather encourage the similarities at the moment than the contrasts.

Duvenhage's paper will appear in *Foundations of Physics*, and I think that is a good thing. Having the paper out in publication-land just encourages someone to write a comment on its deficiencies. Indeed, I would love for that to go further and get a discussion stirring. For, what might get more people to pay attention to the potential of a Bayesian approach to QM than a good stirring discussion? Moreover, if people start to pay attention to the potential of the approach, they might just get to work on filling its other (I would say, more real) deficiencies!

So, there. I can even imagine an eloquent writer to start the ball rolling ... (OK, a hint: His initials are HMW.)

Now, let me go back to the one thing I said "I certainly agree with."

Wisemanism 5: *But here is the difference between classical and quantal.*

In the classical theory the belief p of the weatherman could be given a precise formulation in terms of a set of classical variables (presumably related to his [brain]), and that belief would actually evolve (for a "true h" etc.) in a deterministic way from $p(h)$ to $p(h|d)$ as he looks out the window. That is, the belief of the weatherman can be treated ON THE SAME BASIS as the objects of that belief (i.e. the physical world). That is not to say it MUST be treated on the same basis, but it CAN be, and there is no NEED to have Bayes thm in the foundations of the theory.

In the quantum theory, our belief cannot be treated on the same basis as the physical world. We use a state matrix to encapsulate our belief about the world: our expectation for the results of a fiducial measurement if you like. In the absence of information gathering, this state matrix

evolves in some well-defined way. But if we try to treat our belief on the same basis, as a function of physical brain variables, we run into the QMP.

I think this pinpoints in a pretty terse way the root of our disagreement. To say what you said is to 1) accept a kind of reductionism that I no longer think is healthy, and 2) to posit a strong faith that something *can* be done that, in actual fact, has *never* been done.

Concerning 1) I apologize for using the phrase "those beliefs are presumably a property of his head," for it evokes an imagery I would rather not have in this discussion. Perhaps it would have been better for me to use the word "possession" rather than "property."

Here's a belief I presently possess: With probability greater than .9, there will be another suicide bombing in Israel in the next month. Try to put that into physical terms. What can the word "suicide" mean in classical mechanical terms (or even quantum mechanical terms)? Or to make it look just a little more physical, here is another belief: With probability greater than .99, I will see at least one car today. But why didn't I say a lump of atoms with this characteristic, that characteristic, and the other characteristic? The point is: The world independent of man does not, and cannot, know that that lump of atoms signifies a car. My beliefs—at their starting points—are always about things denoted at such a level of *practical* existence.

Thus I would claim, concerning 2), that it is nothing but a religious faith to suppose that one can derive the form of Bayes' conditionalization rule from a mechanistic physics. It has never been done before, and I would venture to guess that it will never be done.

But still, let me suppose that it could be done after all—it's just that it hasn't been done yet. What would that do? In agreement with (a small part of) Duvenhage's paper, I would say we're left with one conclusion: We should call the difficulties so found the "classical measurement problem." For in classical physics we wouldn't know where the observer begins and the world ends ... just as we don't presently know it with quantum mechanics. Nor as I said above, would we know how to get a Bayesian collapse in the beliefs of the observer (even once he has been identified).

So, with regards to this particular aspect of things, I hold firm with my assessment of Duvenhage. If there is a distinction between a classical and quantum conception of nature, it is somewhere else than in the updating of one's beliefs. I claim the same difficulties are either absent from both conceptions, or present in both conceptions.

But ... maybe I say this all for nothing. For you also wrote this:

Wisemanism 6 : *Duvenhage is quite specific about what he is saying. He says "the Heisenberg cut is therefore no more problematic in quantum mechanics than in classical mechanics". This is true if you are a committed Bayesian, because then the cut is necessarily there from the beginning of your conception of the world.*

and that seems to express that my previous quotation of you was directed more toward Mr. D than me.

In any case,

Wisemanism 7 : *I accept that. This is your solution to the QMP. You say "there is no QMP", but isn't that the same as their saying "I've solved the QMP"? You don't have to answer that.*

I would say it doesn't solve anything; it just shifts the terms. For, where it dismisses old problems, it creates new problems. And it's the problems that keep us all young.

31 May 2002, to Donald, " "Reality" "

Thanks for the clarification on your form of idealism. I also had a nice time reading the FAQ on your webpage. I'm whiling away a little time in Australia at the moment, visiting Nielsen and trying to write some papers with Caves and Schack.

Donaldism 1: *If the "current wavefunction" is just the best description of our knowledge of the system, then what are we made of?*

Something that is not the wavefunction, but for which, once we have accepted its existence, compels us to a structure of reasoning and belief revision whose form is identical to what we once thought of as "quantum mechanics." My best shots so far at saying this in a clear way can be found in Sections 4.2 and 6 of the new paper `quant-ph/0205039`. Or you can see the same thing in pictorial form by looking up the talk "Where's a Good Weatherman When You Need One?" at my webpage.

And so the sparring match goes on ...

01 June 2002, to Grangier, "High Dispute"

I hope you will forgive me for not replying to your letters for so long. The volume of email I have been getting lately has started to become something I am not equipped for. Sometimes I just crack under the pressure, and thus my silence. But at the moment, I've got a little time away from home: I am in Australia, visiting Nielsen and trying to write some papers with Caves and Schack. Beyond that, my wife, kids, and Bell Labs have all been left behind! So I am hoping for a productive three weeks.

Grangierisme 7: *If we speak about "objectivity", we have first to agree about what it is. To keep it simple, I stick with the "naive" view that if I do (or if a student does) a measurement in the lab, this is an objective process: this "did happen", and the fact that you don't know the result of the measurement will not change it.*

As far as I can tell, we do NOT disagree on this. So, it kind of annoys me that you keep bringing it up. Instead, I would say our disagreement lies right here: For some reason YOU think it is NECESSARY to uphold the idea that a pure quantum state is an objective property (of something in nature) in order for your sentence above to come about. Whereas I do not.

I know that I can function just fine with subjective quantum states, even with subjective pure quantum states. You will find no logical inconsistency in me; and I doubt I can find any in you. At the level of our squabble, it is to a large extent a matter of taste. However, you and I both know that matters of taste can lead to matters of fact with regards to the questions we will seek of nature. Your taste will lead you one way; my taste will lead me another. Only history will tell which of us will have had the more productive view. Only history will tell which one of us ended up asking the most interesting questions of nature. (Recently I tried to capture this in a little story to Charlie Bennett and John Smolin about the pleasures of broccoli. I will paste the story at the end of this note.)

Here is the way I put the whole point in a message to Matthew Donald just a minute ago. At the present time, I just do not know how to say it any more cleanly than this:

> If the "current wavefunction" is just the best description of our knowledge of the system, then what are we made of?
>
> Something that is not the wavefunction, but for which, once we have accepted its existence, compels us to a structure of reasoning and belief revision whose form is identical to what we once thought of as "quantum mechanics." My best shots so far at saying this in a clear way can be found in Sections 4.2 and 6 of the new paper `quant-ph/0205039`. Or you can see the same thing in pictorial form by looking up the talk "Where's a Good Weatherman When You Need One?" at my webpage.

The sections I recommended to him, in their focus on the Bureau International des Poids et Mesures, have direct implication on you. So I have a secret dream that you will read them ... and finally, finally something will click in your head, and you will say, "You know, Chris is not being so unreasonable after all."

Grangierisme 8 : *it is quite enough that a complete (=self-defining) set of physical properties can be predicted with certainty. The crucial point is that there is no "ignorance" left (you must admit that each time you write that a pure state has zero entropy).*

It is fine to note that some observables can be predicted with certainty when a state is a pure state, but that does not delete the fact that there is plenty of ignorance left for other observables. One can even quantify that ignorance nicely. See equations (77) through (82) in my `quant-ph/0205039`. (But the result goes back to Wootters.) The von Neumann entropy—which is what you are thinking of in your statement—simply captures the BEST CASE predictability. So what?

Grangierisme 9 : *PS In your paper you strongly "recommend" the work by Lucien Hardy. I have two comments :*
it surprises me that you adhere with the "relative frequency" approach to probabilities that is used by Lucien (his first axiom). I certainly agree with it, but I thought you would not.
Lucien is trying to make QM look like a new probability theory.

If you look at my paper, you will note that I said:

> Beyond that, let me recommend four other articles. The first two are the most technically important for the enterprise I promote in my other contribution to this volume: Namely, to secure a transfer from our present abstract, axiomatic formulation of quantum mechanics to a more physically meaningful one. I think some elements in Lucien Hardy's papers *almost* carry us to the brink of that. In his work, I think the right emphasis is finally being placed on the right mathematical structures. ...
>
> I should point out, however, that in all four of the above references, I think significant improvements could be made by adopting a sufficiently Bayesian stance toward the use and meaning of probability.

I said what I meant. (Just as the experimentalist should strive to perform his measurements accurately; the good theorist should strive to read his friends' papers carefully.) The interesting part of Hardy's papers, in my view, is the mathematics. Saying that constitutes an endorsement of neither 1) the relative-frequency interpretation of probability, nor of 2) Hardy's desire to generalize probability theory to a larger structure. Hardy and I, indeed, have had extensive discussions on this.

Grangierisme 10: *Let us simply admit that local realism is dead, but that physical realism can do quite well without it. Why is that so difficult to accept?*

One more time: I do not deny realism. Moreover, just as you, I am happy with the death of local realism (if what one means by that is "hidden variable realism"). However, I simply do not find your proposed solution to the whole shebang of quantum interpretation problems to be as compelling as you find it.

Measurement clicks alone do not specify post-measurement quantum states. Rather, measurement clicks PLUS prior quantum states (for one's description of the measurement apparatus) do. Thus, where you think a measurement outcome prepares a unique quantum state, I say it has a subjective component. But you say, my measurement device is calibrated. And I say with respect to what? And on and on we could go ad infinitum. The Gordian knot of the state's subjectivity simply cannot be cut by your assuring me that I ought to think otherwise.

Grangierisme 11: *admit that there is a "reality" attached to the pair of particles, but that there is no "reality" attached to each particle.*

This I see as an arbitrary move. Just as one man's unipartite system is half of another man's bipartite system, one man's bipartite system is half of another man's quadripartite system. And so it goes.

Instead, I would say ALL systems have a reality attached to them. It is simply that that reality has NOTHING TO DO with the quantum states one ascribes to those systems. A system has a Hilbert space, and a Hilbert space has a dimension that does not depend upon which state is "alive" within it. That dimension, I would offer you, can be treated as a reality for a quantum system. But there I stop, whereas you want to go further (by supposing a reality to some nonlocal quantum state).

Grangierisme 12: *the "reality" attached to the pair makes no problem with Lorentz invariance, because it was created when the two particles interacted, and it simply follows them if they move very far away. The same conclusion apply to more fancy schemes like entanglement swapping, that requires classical communications to effectively prepare the remote entangled state.*

In the usual sense of what one means by the word "interacted," this—I would say—is just wrong. Suppose Alice and Bob possess two particles that interacted in the past so that they are now entangled (by your way of speaking). Further suppose Carol and Ted possess two particles that interacted in the past so that they are entangled. Maybe even make both states—the AB state and the CT state—to be singlets.

Suppose finally that Alice and Bob have never before seen Carol and Ted. A conclusion one can draw from this is that the A and C particles have never interacted in the past. And the B and T particles have never interacted in the past either.

Let now Alice and Carol meet by chance and perform a Bell measurement on their two particles. If Alice shares all her knowledge of the original AB interaction with Carol, and Carol shares all her knowledge of the original CT interaction with Alice, then they will both be warranted to update their assessments of the BT system. They each will immediately ascribe a pure Bell state (one of four possibilities to that system). With RESPECT to Alice and Carol, the BT system will now be in a Bell state, even though B and T have never interacted in the past.

And this has nothing to do with Alice and Carol sending the information about their measurement outcomes to Bob and Ted. I'll say it one last time: WITH RESPECT TO Alice and Carol, the BT system will be in a Bell state after the measurement.

You see, you get hung up because you want to think of entanglement as a real objective property of two systems, not merely a property of an observer's judgment about those systems. Beyond that, your desire for Lorentz invariance makes you want to think of entanglement as only a consequence of local interaction. But that just goes too far, as the above example shows. Entanglement, just like quantum states, can arise from measurements in the distance.

You can call that a "contextually objective" affair if you wish, i.e., that the entanglement between B and T arises only with respect to the context set by A and C. But then I say—as I've said before—why not just call the state ascribed to the BT system Alice and Carol's information about it? Moreover, by calling it information, you will find that you will stop forgetting the prior (subjective) information that was crucial for defining the posteriori states in the first place.

By the way, I have read your FAQ posted on `quant-ph`; I have not ignored it.

Grangierisme 13 : *Bell's inequalities do not hold !!!! ... if a measurement if performed on one side, ABSOLUTELY NOTHING happens on the other side.*

Well, at least we can agree on this. But look how many of your words I had to delete from the paragraph below to get us there:

Grangierisme 14 : *- since there is no "reality" attached to each particle, Bell's inequalities do not hold !!!! All the job is done by the fact that the individual particles have no quantum state, or no other property whatsoever that would decide on the result of the measurement. Then "action at a distance" simply vanishes: if a measurement if performed on one side, ABSOLUTELY NOTHING happens on the other side.*

And we most certainly agree on the following:

Grangierisme 15 : *What next ? QM is a fantastic theory that can only stimulate one question : why is it working so well ? Then we may notice that QM was invented 75 years ago in a somehow anarchic way, as an attempt to understand atomic spectra. But we may speculate the following : QM is actually the answer to a question that was never clearly formulated. We have the answer, what about finding the question ?*

Even if we do not see eye to eye in the secular world, we seem to dream of the same heaven.
PS. Don't forget the story below.

King Broccoli

"In any case, none of this matters for doing physics," you say. But I think it does in the long run. (Certainly, you've got to concede that there's something that fuels me—and it hurts to think that you might think it is nothing more than irrationality.) When Charlie sent me the picture of a skunk cabbage in reply to these very issues, I found myself thinking of King Broccoli. The story goes that one day, by divine providence, it came to King Broccoli that broccoli, the vegetable, his namesake, actually tastes good. Good in a way that hitherto only gods and angels had known. Every child who had ever said, "Yuck, I don't like broccoli; it tastes awful!" was simply wrong ... or at least that's what the king realized. King Broccoli, being the head of state, decided to do something about it. Henceforward, all gardens in the kingdom should have a patch devoted solely to broccoli. It really wasn't much of a burden on the national product (except, perhaps, for the psychiatrists who had to treat all the movie stars who had never felt fulfilled in their broccoli experience). But think of the diversity of vegetables

the kingdom might have raised if its citizens hadn't been encumbered with the king's notion that broccoli had an objective, but never verifiable, taste?

Moral? Maybe there is none. But it is a documented fact that the Kingdom of Broccoli eventually fell and was replaced by a liberal democracy (where the ideals set forth in the constitution, rather than the particulars laws of any given day, have an objective status).

12 June 2002, to Donald, "Receipt"

I got your long note; thanks. I'm going to mull over it for a while before replying. But I will reply.

Donaldism 2: *I also haven't commented much on the points I agree with in the paper, although there are a few of these!*

It would be nice to see what these are. In ways, I'm more interested in where we agree than in where we disagree. A good discussion needs some solid ground somewhere.

Beyond that, it would please me to no end to learn that you might have found a thought or two that you could use in the thing. But if you didn't, you didn't. It's a horrible feeling to think I might be writing nothing for nothing.

Each day, I tell my daughters that they can change the world. I tell them that they can change it to the core. But I never tell them that they can believe anything they want. There is a difference. And you don't see it.

The summary at the end of your note troubles me to no end.

Donaldism 3: *Your know-nothing ism, like de Finetti's irrationalism (Gillies, "Philosophical Theories of Probability", page 86), have the dangers of Bohr's writings on which I would agree with Beller (Physics Today, September 1998, pages 29 - 34). In particular, by leaving far too much in vagueness, incoherence, and pious hope, you give the religously-minded the official endorsement of the physics establishment that they may believe anything they want, instead of, by example, instructing them that they can believe anything they want as long as it is rational, coherent, tentative, revisable, and compatible with the evidence (and therefore contrary to naive expectations, because if quantum theory, or indeed science in general, tells us anything it is that the world is not how we would have imagined it before we investigated); and they accept that they may be completely wrong.*

I will certainly return.

24 June 2002, to Wiseman, "Points"

You see, I'm catching up on my email finally. Thanks for the Nash pointer and the Aussie WWII statistics.

Wisemanism 8: *4. re. the processing power limitations on quantum states. You could have a look at my paper "Bayesian feedback versus Markovian feedback in a two-level atom" (Phys Rev A 2002, or* `quant-ph`*) to see what language I actually did use (I can't remember). But the point is that the state an experimenter computes on the fly to do feedback control is not necessarily less pure than the "true" state which the experimenter does not know. But the experimenter still wants to know how much his state is likely to be "in error" from the "true state". I guess so the experimenter armed*

with this knowledge would have a better state estimate, by accounting for the error by making the state more mixed. But in any case, how can you call the "true state" a state of belief or whatever if there is no one who believes it? It is what the experimenter's state of knowledge should be, given what they know. That is, it is their state of knowledge, even if they are too dumb to know it. This is another reason you may use to avoid the phrase "state of knowledge", but the challenge is still how to describe this "true state" in that case, without circumlocuting intolerably.

What you are asking for is the experimenter to compare his pragmatic gambling commitments under the real world constraints he is living under at a given moment, to the beliefs or commitments he would possess under more ideal conditions (for example, if he had limitless computing power). Just as there is no "true" quantum state, independent of the agent—i.e., for two agents, there might be two quantum states—the same may be the case for a single agent. I like your examples because they help draw that out. I can contemplate how I should bet believing what I believe now and knowing that my feeble mind cannot analyze the full implications of those beliefs, or I can contemplate how I should bet believing what I believe now and imagining a supercomputer that draws out the proper conclusions from those beliefs.

In all cases, whatever the final products of all calculations and all approximations one writes down, those are one's gambling commitments. And *with respect to* that situation, that is one's "quantum state" for a system.

24 June 2002, to Wiseman, "The World is Under Construction"

Wisemanism 9: *Do you believe that events in the world really are random? Or do you believe they only appear to be random? In the first case, doesn't that mean that you have to believe in objective probabilities? . . .*

Or are you saying that the real world is unanalysable, unthinkable even? Everything we say should be couched in terms of gambling commitments. First, that seems to be a cop-out, giving up on any understanding of the Universe. Second, it can't explain anything in the Darwinian way you mentioned, except Dutch-book consistency. It can't explain why it is "bad" to hold a gambling commitment based on the idea that all world cup soccer balls contain bombs that have a 50% chance to blow up every time a goal is scored . . . You cannot say anything about animals that would have been likely to have gone extinct because of poor (but consistent) gambling commitments, because that is a statement using the concept of objective probabilities. You cannot explain *anything that is not strictly deterministic without using objective probabilities, it seems to me.*

I trust you understand my motives. I wouldn't bother discussing this with you if I didn't think your ideas were potentially revelational. What does not kill you makes you stronger.

For a couple of days I have been thinking about how to reply to the questions in your 6/14/02 email 'reality', but this morning I found myself significantly revising the response I had started to build. In particular, I decided to hardly reply to your questions at all! This may be a little bit annoying to you, but I think it will benefit our longer term discussion.

Of the three options you gave me for answering your questions (I only quoted the last one of the three above), I suppose if I were forced to choose one, I would align myself with the one you called a "cop-out." However, from my point of view, the language you use builds about the ugliest picture it can for where this effort is going. Indeed, you miss the very point, the very beauty, of the "cop-out." So, what I'd like to do is set that right—right here and right now—before we go much further. As advertised, in that way, I will not reply straightforwardly to your questions.

You see, the very starting point for most of my latest thoughts—the thing I think quantum mechanics gives us the deepest and most thorough hint of—is that there is no such thing as THE universe in any completed and waiting-to-be-discovered sense. The thought I am *testing out* is that the universe as a whole is still under construction. And when I say this, I am not thinking of just bits and pieces of it; I am thinking of the whole shebang, all the way to the roots. Nothing is completed. Not just the playhouse Kiki is building for Emma and Katie, or the evolutionary track of the human species, but even the "very laws" of physics. The idea is that they too are building up in precisely the way—and ever in the same danger of falling down as—individual organic species. That is to say, it's Darwinism all the way down.

So when you ask me if I am "saying that the real world is unanalysable, unthinkable even," the answer in a way is "yes." For it is blatantly impossible to analyze to the last detail the characteristics of a world that has not even been dreamt up (even in its own mind's eye).

But how can I impress this upon you, or even make it seem reasonable as a direction for research? That is a tough call. For, like with beer or single-malt Scotch, it is surely an acquired taste that builds only slowly and with the right company. Of course, I could just send you back to my paper `quant-ph/0204146` and ask you to take it very seriously. But this morning it dawned on me to maybe spend a little time with my scanner to try to IV some thoughts straight into your bloodstream.

At the moment, I can think of no better introductions to the line of thought I'd like to expose you to than three articles by Richard Rorty: "A World without Substances or Essences," "Truth without Correspondence to Reality," and "Thomas Kuhn, Rocks, and the Laws of Physics."[15] (Read them in that order, if you read them.) All three papers can be found in his collection of essays, *Philosophy and Social Hope*. If you absorb these, I think you'll understand completely what I'm up to, and why I so dislike the negative connotations you associate with the radical-Bayesian way of viewing the quantum state. Of course, it may not turn your head the way it turns mine, but at least you'll know where I'm coming from, and from what pool of enthusiasm I derive my strength to eschew the "golden nuggets" of *mere* quantum cosmology, *mere* Bohmianism, and *mere* "dreams of a final theory." The world as I see it is a much bigger place than those stories can tell. And the interpretational issues at the core of quantum mechanics strike me as our first rigorous indication that there is something more to this idea than simply the hopes and desires of an enthusiast.

For now, let me give you a flavor of the thoughts in these papers, and then leave you on your own in the case that you would like to pursue this further. The following quotes come from "Truth without Correspondence to Reality."

> In this essay I shall focus on Whitman's phrase 'counts ... for her justification and success ... almost entirely upon the future'. As I see it, the link between Whitmanesque Americanism and pragmatist philosophy—both classical and 'neo-'—is a willingness to refer all questions of ultimate justification to the future, to the substance of things hoped for. If there is anything distinctive about pragmatism it is that it substitutes the notion of a better human future for the notions of 'reality', 'reason' and 'nature'. One may say of pragmatism what Novalis said of Romanticism, that it is 'the apotheosis of the future'.

[15] WARNING: Just because I say I can think of no better introductions to these ideas, it does not mean I endorse every statement in these papers; I may not endorse even half of them. However, I think these papers go in the right direction, even if they go too far ... and even if their arguments are far too weak. But I choose the papers I do because they are easy reading, with beautiful writing, and I suspect these thoughts are so foreign to you that if you can find any sense in *some of them*, then it may be a good start for a dialogue. Moreover, I continue to stress that the best justification yet to pursue this direction of thought—and this is something Rorty does not know—is quantum mechanics itself. So, rather than being the final words on things, these are just the beginning words on things.

As I read Dewey, what he somewhat awkwardly called 'a new metaphysic of man's relation to nature', was a generalization of the moral of Darwinian biology. The only justification of a mutation, biological or cultural, is its contribution to the existence of a more complex and interesting species somewhere in the future. Justiflcation is always justification from the point of view of the survivors, the victors; there is no point of view more exalted than theirs to assume. This is the truth in the ideas that might makes right and that justice is the interest of the stronger. But these ideas are misleading when they are construed metaphysically, as an assertion that the present status quo, or the victorious side in some current war, stand in some privileged relation to the way things really are. So 'metaphysic' was an unfortunate word to use in describing this generalized Darwinism which is democracy. For that word is associated with an attempt to replace appearance by reality.

Pragmatists—both classical and 'neo-'—do not believe that there is a way things really are. So they want to replace the appearance-reality distinction by that between descriptions of the world and of ourselves which are less useful and those which are more useful. When the question 'useful for what?' is pressed, they have nothing to say except 'useful to create a better future'. When they are asked, 'Better by what criterion?', they have no detailed answer, any more than the first mammals could specify in what respects they were better than the dying dinosaurs. Pragmatists can only say something as vague as: Better in the sense of containing more of what we consider good and less of what we consider bad. When asked, 'And what exactly do you consider good?', pragmatists can only say, with Whitman, 'variety and freedom', or, with Dewey, 'growth'. 'Growth itself,' Dewey said, 'is the only moral end.'

They are limited to such fuzzy and unhelpful answers because what they hope is not that the future will conform to a plan, will fulfil an immanent teleology, but rather that the future will astonish and exhilarate. Just as fans of the avant garde go to art galleries wanting to be astonished rather than hoping to have any particular expectation fulfilled, so the finite and anthropomorphic deity celebrated by James, and later by A. N. Whitehead and Charles Hartshorne, hopes to be surprised and delighted by the latest product of evolution, both biological and cultural. Asking for pragmatism's blueprint of the future is like asking Whitman to sketch what lies at the end of that illimitable democratic vista. The vista, not the endpoint, matters.

So if Whitman and Dewey have anything interesting in common, it is their principled and deliberate fuzziness. For principled fuzziness is the American way of doing what Heidegger called 'getting beyond metaphysics'. As Heidegger uses it, 'metaphysics' is the search for something clear and distinct, something fully present. That means something that does not trail off into an indefinite future ...

and

So far I have been trying to give an overview of Dewey's place in the intellectual scheme of things by saying something about his relation to Emerson, Whitman, Kant, Hegel and Marx. Now I want to become a bit more technical, and to offer an interpretation of the most famous pragmatist doctrine—the pragmatist theory of truth. I want to show how this doctrine fits into a more general programme: that of replacing Greek and Kantian dualisms between permanent structure and transitory content with the distinction between the past and the future. I shall try to show how the things which James and Dewey said about truth were a way of replacing the task of justifying past

custom and tradition by reference to unchanging structure with the task of replacing an unsatisfactory present with a more satisfactory future, thus replacing certainty with hope. This replacement would, they thought, amount to Americanizing philosophy. For they agreed with Whitman that America is the country which counts for its 'reason and justification' upon the future, and *only* upon the future.

Truth is what is supposed to distinguish knowledge from well-grounded opinion—from justified belief. But if the true is, as James said, 'the name of whatever proves itself to be good in the way of belief, and good, too, for definite, assignable, reasons', then it is not clear in what respects a true belief is supposed to differ from one which is merely justified. So pragmatists are often said to confuse truth, which is absolute and eternal, with justification, which is transitory because relative to an audience.

Pragmatists have responded to this criticism in two principal ways. Some, like Peirce, James and Putnam, have said that we can retain an absolute sense of 'true' by identifying it with 'justification in the ideal situation'—the situation which Peirce called 'the end of inquiry'. Others, like Dewey (and, I have argued, Davidson), have suggested that there is little to be said about truth, and that philosophers should explicitly and self-consciously *confine* themselves to justification, to what Dewey called 'warranted assertibility'.

I prefer the latter strategy. Despite the efforts of Putnam and Habermas to clarify the notion of 'ideal epistemic situation', that notion seems to me no more useful than that of 'correspondence to reality', or any of the other notions which philosophers have used to provide an interesting gloss on the word 'true'. Furthermore, I think that any 'absoluteness' which is supposedly ensured by appeal to such notions is equally well ensured if, with Davidson, we insist that human belief cannot swing free of the nonhuman environment and that, as Davidson insists, most of our beliefs (most of *anybody's* beliefs) must be true. For this insistence gives us everything we wanted to get from 'realism' without invoking the slogan that 'the real and the true are "independent of our beliefs" '—a slogan which, Davidson rightly says, it is futile either to accept or to reject.

Davidson's claim that a truth theory for a natural language is nothing more or less than an empirical explanation of the causal relations which hold between features of the environment and the holding true of sentences, seems to me all the guarantee we need that we are, always and everywhere, 'in touch with the world'. If we have such a guarantee, then we have all the insurance we need against 'relativism' and 'arbitrariness'. For Davidson tells us that we can never be more arbitrary than the world lets us be. So even if there is no Way the World Is, even if there is no such thing as 'the intrinsic nature of reality', there are still causal pressures. These pressures will be described in different ways at different times and for different purposes, but they are pressures none the less.

The claim that 'pragmatism is unable to account for the absoluteness of truth' confuses two demands: the demand that we explain the relation between the world and our claims to have true beliefs and the specifically epistemological demand either for present certainty or for a path guaranteed to lead to certainty, if only in the infinitely distant future. The first demand is traditionally met by saying that our beliefs are made true by the world, and that they correspond to the way things are. Davidson denies both claims. He and Dewey agree that we should give up the idea that knowledge is an attempt to *represent* reality. Rather, we should view inquiry as a way of using reality. So the relation between our truth claims and the rest of the world is causal rather than

representational. It causes us to hold beliefs, and we continue to hold the beliefs which prove to be reliable guides to getting what we want. Goodman is right to say that there is no one Way the World Is, and so no one way it is to be accurately represented. But there are lots of ways to act so as to realize human hopes of happiness. The attainment of such happiness is not something distinct from the attainment of justified belief; rather, the latter is a special case of the former.

Pragmatists realize that this way of thinking about knowledge and truth makes certainty unlikely. But they think that the quest for certainty—even as a long-term goal—is an attempt to escape from the world. So they interpret the usual hostile reactions to their treatment of truth as an expression of resentment, resentment at being deprived of something which earlier philosophers had mistakenly promised. Dewey urges that the quest for certainty be replaced with the demand for imagination—that philosophy should stop trying to provide reassurance and instead encourage what Emerson called 'self-reliance'. To encourage self-reliance, in this sense, is to encourage the willingness to turn one's back both on the past and on the attempt of 'the classical philosophy of Europe' to ground the past in the eternal. It is to attempt Emersonian self-creation on a communal scale. To say that one should replace knowledge by hope is to say much the same thing: that one should stop worrying about whether what one believes is well grounded and start worrying about whether one has been imaginative enough to think up interesting alternatives to one's present beliefs. As West says, 'For Emerson, the goal of activity is not simply domination, but also provocation; the telos of movement and flux is not solely mastery, but also stimulation.'

and

It may seem strange to say that there is no connection between justification and truth. This is because we are inclined to say that truth is the aim of inquiry. But I think we pragmatists must grasp the nettle and say that this claim is either empty or false. Inquiry and justification have lots of mutual aims, but they do not have an overarching aim called truth. Inquiry and justification are activities we language-users cannot help engaging in; we do not need a goal called 'truth' to help us do so, any more than our digestive organs need a goal called health to set them to work. Language-users can no more help justifying their beliefs and desires to one another than stomachs can help grinding up foodstuff. The agenda for our digestive organs is set by the particular foodstuffs being processed, and the agenda for our justifying activity is provided by the diverse beliefs and desires we encounter in our fellow language-users. There would only be a 'higher' aim of inquiry called 'truth' if there were such a thing as *ultimate* justification—justification before God, or before the tribunal of reason, as opposed to any merely finite human audience.

But, given a Darwinian picture of the world, there can be no such tribunal. For such a tribunal would have to envisage all the alternatives to a given belief, and know everything that was relevant to criticism of every such alternative. Such a tribunal would have to have what Putnam calls a 'God's eye view'—a view which took in not only every feature of the world as described in a given set of terms, but that feature under every other possible description as well. For if it did not, there would remain the possibility that it was as fallible as the tribunal which sat in judgment on Galileo, a tribunal which we condemn for having required justification of new beliefs in old terms. If Darwin is right, we can no more make sense of the idea of such a tribunal than we can make sense of the idea that biological evolution has an aim. Biological evolution

> produces ever new species, and cultural evolution produces ever new audiences, but there is no such thing as the species which evolution has in view, nor any such thing as the 'aim of inquiry'.
>
> To sum up, my reply to the claim that pragmatists confuse truth and justification is to turn this charge against those who make it. They are the ones who are confused, because they think of truth as something towards which we are moving, something we get closer to the more justification we have. By contrast, pragmatists think that there are a lot of detailed things to be said about justification to any given audience, but nothing to be said about justification in general. That is why there is nothing general to be said about the nature or limits of human knowledge, nor anything to be said about a connection between justification and truth. There is nothing to be said on the latter subject not because truth is atemporal and justification temporal, but because *the* **only** *point in contrasting the true with the merely justified is to contrast a possible future with the actual present.*

I don't have to tell you that I find these ideas tremendously exciting. It is not that nature is hidden from us. It is that it is not all there yet and never will be; 'nature' is being hammered out as we speak. And just like with a good democracy, we all have a nonnegligible input into giving it shape. That is the idea I am *testing* for consistency and utility. On the chance that it will lead somewhere, it seems to me, worth the gamble.

26 June 2002, to Pitowsky, "The One Boolean Algebra Approach"

Let me make one further comment. I too used to think that the PBA approach was the way to go if one wanted to build up a theory along quantum logical lines. But now, I'm not so convinced of it. That is because I am starting to think that quantum mechanics is more analogous to the epistemological theory Richard Jeffrey calls "radical probabilism" than anything else. From that view, there are "probabilities all the way down" with one never getting hold of the truth values of *any* propositions. Rüdiger Schack and I just discovered a wealth of material on Jeffrey's webpage `http://www.princeton.edu/~bayesway/`.

In any case, I think what this leads to is that we ought to be focusing much more on characterizing quantum mechanics solely in terms of the "logic" of POVMs than anything else—these being the structures analogous to what crops up in Jeffrey's "probability kinematics." Thus, if one is looking to characterize PBAs, the best task might be to focus on the kinds of PBAs that POVMs generate, rather than the ones of Kochen and Specker based solely on standard measurements. (This may or may not have some connection to what people are calling "effect algebras" but I don't know.)

Beyond that, I am now of the mind that all one really ever needs for understanding quantum mechanics is a *single* Boolean algebra that is kept safely in the background (solely) for reference. The rest of the theory (and indeed all real-world measurement) is about probability kinematics with respect to that algebra. See Sections 4.2, 6, and 6.1 of my paper `quant-ph/0205039` for details.

27 June 2002, to Poulin, "Compatibility Never Ends"

I apologize for taking so long to get back to you. My email box has just been running over lately, and on top of that I've been traveling a lot (just back from Australia, actually).

Poulinism 1: *About your approach: as far as I can see, the main difference (aside from the "vocabulary") between your derivation of the BFM compatibility criterion and their's is that you are aware that it follows from* strong consistency*.*

No, I think the difference between us and them runs deeper than that. Let me try to express the point in a way that maybe I haven't used before. Concerning the BFM criterion: When BFM call it "necessary and sufficient" and when CFS call it "necessary and sufficient," we mean two very different things. For BFM the "necessary" part is enforced by the supposed existence of a super-observer, Zeno. That is, they mean that their criterion is "necessary" because they suppose a super-observer, Zeno, must exist. (See the final section of their paper.) What we (i.e., CFS) do instead is define a notion of "compatibility" based on the classical notion of "strong consistency" and show that the BFM condition is necessary and sufficient with respect to that. Beyond that, however, we leave the issue open.

That is, with regard to the question, "MUST the quantum states of two observers be compatible?," we make no statement in the paper. That is because, from our view—or at least mine in particular—there is nothing in nature that enforces that states MUST BE compatible. This is quite important if one wants to get a consistent Bayesian view off the ground. (I try to say all this in a lengthier and maybe more complete way in my samizdat *Quantum States: What the Hell Are They?*, pages 116–118, in a note titled "The Spirit of Gandhi." Actually, I think I say it better there; so have a look at that.)

Now, if there is nothing to enforce compatibility in the BFM sense, what would happen if two observers are incompatible with respect to this criterion? That's a fair and decent question, and I think where all the excitement begins. What the W criterion shows is that the two observers can just shake off their incompatibility and go forward (if they are in a congenial mood). That is, they can come to future agreement, not simply by pooling their prior beliefs, but by making an active intervention upon the system their states are about (i.e., by making a measurement). I view this as an extremely interesting property of the quantum world: it allows possibilities for going forward that the classical world does not.

On the other hand, if two observers are not in a congenial mood (i.e., will not perform a measurement of the W variety), then a "crisis" can ensue. What the full implications of that are, we don't know yet. But it may be even more exciting still.

Poulinism 2: *Maybe honesty can be related to strong consistency?*

Yes, you are right, and this is an important point. I discuss this at length in a note titled "The Commitments" on pages 133–138, and again on pages 192–193 in a note titled "More Balking." The upshot is, yes, strong consistency is deeply tied to honesty. The next question is, must one enforce honesty in a Bayesian approach to things? Your note implies that you tend to think, "yes." The pages I refer you to show that I tend to think, "no."

Poulinism 3: *Finally, while you relate the BFM condition to your ES, PP, W, and W', it is fundamentally different in spirit. This distinction is related to what we refer to as "type of knowledge". In our vocabulary, the BFM criterion is appropriate to compare knowledge of the quantum while the other are suited for quantum knowledge (just like quantum fidelity is). To me, the most convincing argument is the one of two nonorthogonal pure states: the "measurement criteria" will grant that they are compatible while, as we expressed it "if two observers claim to have complete knowledge of a system, their descriptions had better agree completely." Comparing states of knowledge and measurement outcomes predicted from these states are two different things. We have discussed a similar subtlety by comparing knowledge described by a density matrix and*

one described by a preparation: these are two different questions which deserve different answers although physically, they cannot be distinguished.

It's a question of perspective. If one goes the full Bayesian route (which I am starting to do), then it is much better to call a quantum state a "state of belief or judgment," and not a "state of knowledge." For then, one finds that one is never inclined to say "their descriptions had better agree completely." There is nothing to enforce that, except possibly Darwinism. (I'm serious about that statement.) That is to say, when the world rears its head, someone with a firm belief might be wiped out, but in a Bayesian approach (where probabilistic statements always depend upon SUBJECTIVE priors), there is nothing to enforce compatibility beforehand. A misjudgment can only be declared a misjudgment after the fact, not before the fact (which in the quantum setting is created by the process of measurement).

27 June 2002, to Wiseman, "Probabilism All the Way Up"

Wisemanism 10: *Second, my wife, Nadine, wants to know why Kiki was building a playhouse while you indulge yourself in philosophy.*

I'd like to think we both do what we do best. But I suspect there's no philosopher out there who would say I'm doing good philosophy. (Kiki's artistry on the other hand is always a good hit.)

Wisemanism 11: *Now, more seriously, you say that my language "builds about the ugliest picture it can for where this effort is going". As I keep saying, I mean to be provocative. I hope it drives you to new heights in building a beautiful picture in response. Honestly I do see the beauty in your program. And I think the more extreme it becomes, the more beautiful it becomes. I am very interested to see where it ends up.*

Thanks for the compliment. And, indeed, your correspondence does drive me to new heights (of something). But now I worry that I offended you with my phrase "ugliest picture." It probably came off that way, but it wasn't meant to be an emotional statement or a point about you personally. If some emotion did slip into it, it most likely refers to a conversation I had with Harvey Brown, circa September 11 of last year. Harvey kept saying that I wanted to "doom" nature to being "ineffable." But that language carries such a negative connotation. It carries the idea that there is something there that we can never, or should never, attempt to speak of. So, when you said something similar in print, it gave me the opportunity to try to reply in print. (As you know, I try to have my thoughts recorded so I can refer people to them. One of the original ideas was that it would save me time that way; so far, that aspect of it hasn't worked out.) Anyway, as I made clear, I want to combat that with all my strength. In particular, the way that I am thinking about it, it is not a bad thing that there are some things beyond description in nature. Instead, it is just a statement that there are more things to come; it is a way of leaving room for something new.

Wisemanism 12: *As it happens, I don't have much of a taste for beer or single-malt Scotch. Also as it happens, I was reading a critique of Richard Rorty the very morning before I got your letter. Otherwise I never would have heard of him. It was a 1997 article by Alex Callinicos "Post-modernism: a critical diagnosis". The most interesting criticism in there was to say that Rorty "presumes what he needs to establish, namely that science and philosophy can be assimilated into literature. ... It is ... very hard in practice when trying to explain why one theory can be said to be more useful than another to avoid at least tacitly appealing to the idea that it captures how things are better than its rival does."*

Perhaps this is one aspect of Rorty you disagree with. But I wonder about your saying that quantum mechanics is the best justification for Rorty's philosophy, as if quantum mechanics is something you accept to be real, an "intrinsic nature of reality", the very idea of which Rorty explicitly rejects.

First, just a technical point. The philosophies I am most attracted to at present are those of James and Dewey and what James says about F.C.S. Schiller (but I haven't read Schiller himself yet). Rorty has donned himself to be the spokesman of those guys—and I don't mind that because he writes so nicely—but his writings also have a good admixture of the postmodernist ideas (of Foucault, Derrida, etc.) thrown into them to boot. This business about science not being more trustworthy or real than literary criticism presently strikes me as going too far.

But to Callinicos' point (give me the reference, by the way, and I'll read it)—"It is ... very hard in practice when trying to explain why one theory can be said to be more useful than another to avoid at least tacitly appealing to the idea that it captures how things are better than its rival does."—I would just reply, "Darwinism." And then, if that didn't sink in, I'd say, "Darwinism." The point is, from this conception, there is very little to say beyond that. Were elephants written into the blueprints of the universe? From the Darwinistic conception, they were not. Yet, the species fills a niche and has had a stability of at least a few million years worth. There is a sense in which an elephant, like a theory, is a "true" component in a description of the world. But that "trueness" only has a finite lifetime, and is largely a result of a conspiracy of things beyond its command (selection pressures). To put it another way, in contrast to Callinicos, the elephant doesn't "capture how things are better than its rival does" in any absolute sense—only in a transitory sense—but that doesn't take away from the functional value of the elephant today. So too, I am trying to imagine with theories.

Henry Folse, by the way, wrote me that there is something of a tradition with this evolutionary idea (beyond Rorty). So I've got a big reading list ahead of me: He tells me Toulmin, Kuhn, Kitcher, and van Fraassen.

Now, to quantum mechanics. You find something contradictory about my liking both quantum mechanics and Rorty. Here is the way I would put it. Presently at least, I am not inclined to accept quantum mechanics "to be real, an 'intrinsic nature of reality'," except insofar as, or to the extent that, it is a "law of thought," much like simple (Bayesian) probability theory. Instead, I view quantum mechanics to be the first *rigorous* hint we have that there might actually be something to James's vision.

I've already told you the history of this, haven't I? I gave a talk in 1999 at Cambridge on the quantum de Finetti theorem, after which Matthew Donald came up to me and bellowed, "You're an American pragmatist!" I didn't know what that meant really, but I kept the thought in the back of my head; I figured one day, I'd figure out what he meant. As it goes, that happened on July 21 of last year. I came across this book of Martin Gardner's of which one of the chapters was titled, "Why I Am Not a Pragmatist." (Part of the story is recorded on page 15 of my little samizdat in a note titled "The Reality of Wives." You might read it for a little laugh.) As I read it, it was like a flash of enlightenment. For every reason Gardner gave for not being a pragmatist, I thought about quantum mechanics and realized that indeed I was one. Donald was right after all; I am an American pragmatist. And my further study of pragmatism has borne that out to a T.

My point of departure, unlike James's, was not abstract philosophy. It was simply trying to make sense of quantum mechanics, where I think the most reasonable and simplest conclusion one can draw from the Kochen-Specker results and the Bell inequality violations is, as Asher Peres says, "unperformed measurements have no outcomes." The measurement provokes the "truth value" into existence; it doesn't exist beforehand. Now, go off and read about James's and Dewey's theory of

truth and you'll find almost exactly the same idea (just without the rigor of quantum mechanics). And similarly with lots of other pieces of the philosophy.

So, I view quantum mechanics as the hint of something much deeper. But the full story is not yet told. That is, quantum mechanics strikes me as being to our community what the Galapagos Islands were to Darwin—just a hint of something bigger.

Wisemanism 13: *You and Rorty I guess would agree that "dreams of a final theory" will never be more than dreams. I guess that idea does not worry me as much as it would some physicists, but it does seem like a defeat. But perhaps that just says something of my personality. How much of a role does personality play in one's preferred philosophy?*

Your question is a good one, and one I worry about a lot. Where your knee-jerk reaction is defeat, mine is one of unlimited possibilities and newfound freedom. On a similar issue, James put it like this:

> The history of philosophy is to a great extant that of a certain clash of human temperaments. Undignified as such a treatment may seem to some of my colleagues, I shall have to take account of this clash and explain a good many of the divergencies of philosophies by it. Of whatever temperament a professional philosopher is, he tries, when philosophizing, to sink the fact of his temperament. Temperament is no conventionally recognized reason, so he urges impersonal reasons only for his conclusions. Yet his temperament really gives him a stronger bias than any of his more strictly objective premises. It loads the evidence for him one way or the other, making a more sentimental or more hard-hearted view of the universe, just as this fact or that principle would. He *trusts* his temperament. Wanting a universe that suits it, he believes in any representation of the universe that does suit it. He feels men of opposite temper to be out of key with the world's character, and in his heart considers them incompetent and 'not in it,' in the philosophic business, even though they may far excel him in dialectical ability.
>
> Yet in the forum he can make no claim, on the bare ground of his temperament, to superior discernment or authority. There arises thus a certain insincerity in our philosophic discussions: the potentest of all our premises is never mentioned. I am sure it would contribute to clearness if in these lectures we should break this rule and mention it, and I accordingly feel free to do so.

But I think the disparity between our views is in better shape than that. I think you're only seeing the program "physics is the ability to win a bet" as a defeat because—even if you don't know it—you're working within a kind of Kantian mindset. That the universe is already formed and there; that there is an "a priori." Anything that can't be said about the universe is then most surely a loss or limitation. But, I think once you see that what the pragmatist is trying to get at is not that, maybe your heart will change. Physics as the ability to win a bet will strike you as something immensely positive. Physics is like that because reality is still forming, and the Darwinistic component (along with the "non-detachedness" of the observer in quantum mechanics) indicates that it may be somewhat malleable. From that point of view, to have "dreams of a final theory" is almost like admitting defeat.

But given what you've said, maybe you're already starting to feel some of this. And that tickles me immensely.

Wisemanism 14: *To conclude, I can (or rather could) accept a lot, or even all, of what you and Rorty are exploring. But I am still not emotionally or intellectually compelled to do so. And I am*

really not sure whether I want to be compelled in one direction, or whether I want to be able to contain conflicting philosophies. I have this idea that there is an incoherence at the heart of things. Irreconcilable levels of description. Profound truths being the opposite of other profound truths. Incompleteness theorems. That sort of stuff.

Understood. I know that there's nothing worse than an evangelist knocking at your door on a Saturday morning. Feel free to not reply to this note at all. In the mean time, I'll try to do my best to do what I really ought to be doing: proving theorems, simplifying the quantum axioms, trying to find real-world physics problems for which this view is the most powerful way to tackle it, etc., etc. One thing physicists never deny is a better method for solving a problem.

28 June 2002, to Poulin, "To Believe to Know"

Poulinism 4: *Otherwise, you get incompatible statements with which Bayesian theory cannot deal (like computing $p(x|y)$ when y is assigned probability 0).*

It's starting to sound like we mean two different things by Bayesian theory. See, `http://www.princeton.edu/~bayesway/KC.tex.pdf`, slide #12 (on page 13 actually). More seriously, see: `http://cepa.newschool.edu/het/essays/uncert/subjective.htm`, especially the parts on Ramsey and de Finetti.

Poulinism 5: *I think that most of the discussion is based around this distinction: "state of knowledge" and "state of belief". I would say that states of knowledge must be (BFM) compatible while states of belief can be incompatible. How are these two things defined? Well I would say that a state of knowledge is built with the help of an initial state on which everybody agrees, the postulates of quantum mechanics, and "public actions".*

I will agree to your definition of "state of knowledge." But, backtracking from that, an initial state upon which everyone agrees? If one is taking a subjectivist approach (or what I had been calling a "Bayesian approach") to interpreting the quantum state, there is nothing in nature to enforce an initial prior agreement. God does not come down from on high and say to all the agents (i.e., all the observers), "Your starting point shall be the quantum state Ψ." Everyone is left to fend for himself.

That is to say, in the language of the second webpage I sent you to above, in a world where the quantum state is not presupposed to be objective—as had been Mermin's goal when he started this exercise—there is nothing to enforce the "Harsanyi doctrine."

28 June 2002, to Caves-Schack, "The Harsanyi Doctrine"

I just happened to run across the following discussion:

> Finally we should mention that one aspect of Keynes's (1921) propositions has re-emerged in modern economics via the so-called "Harsanyi Doctrine"—also known as the "common prior" assumption (e.g. Harsanyi, 1968). Effectively, this states that if agents all have the same knowledge, then they ought to have the same subjective probability assignments. This assertion, of course, is nowhere implied in subjective probability theory of either the Ramsey-de Finetti or intuitionist camps. The Harsanyi doctrine is largely an outcome of information theory and lies in the background of

rational expectations theory—both of which have a rather ambiguous relationship with uncertainty theory anyway. For obvious reasons, information theory cannot embrace subjective probability too closely: its entire purpose is, after all, to set out a objective, deterministic relationship between "information" or "knowledge" and agents' choices. This makes it necessary to filter out the personal peculiarities which are permitted in subjective probability theory.

29 June 2002, to Poulin, "Incompatible Beginnings"

Poulinism 6: *Maybe requiring a single description on which everyone agrees is a bit too strong, but something has to be imposed.*

The whole point of the research program I am building up, and the point of most of the 220 pages in *Quantum States: What the Hell Are They?*, and much of the point of my paper `quant-ph/0205039`, is that I don't buy that. Remember the line in the Pink Floyd song, "Mother's going to put all her fears into you."? The point is, none of us are born with "rational" views. We start out in life with whatever our community pumps into us. The mark of "rationality" instead is how we change our views in the light of evidence, and how we gamble based on what we believe. That is what Dutch-book coherence is about, for instance.

So, when I say "agents are left to fend for themselves," I mean it. Every now and then agents start out compatible in one sense or another and then they stay that way, indeed ... perhaps with ever more similarities between their distributions as data flows in. But that is not the norm. Instead most of the time, there is some incompatibility in our world views lurking in the background. We can live with that as long as nothing comes to light to challenge our views. But when the opportunity arises, there will be a "crisis," as Ruediger Schack likes to say.

At that point, as you say, Bayesian theory can't handle things. But the view I am shooting for is that Darwinism can. That is, if you have a firm belief (i.e., a probability 1 ascription) for something, by Dutch-book rules you are willing to bet your whole bank account on the event. If the event doesn't occur, then you have lost your whole bank account, and in a way your life.

The program, as I see it, is to see how rigorous and fruitful on can make that line of thought.