The Bare Theory Has No Clothes

Jeffrey Bub, Rob Clifton, and Bradley Monton

MINNESOTA STUDIES IN THE PHILOSOPHY OF SCIENCE

VOLUME XVII

Quantum Measurement: Beyond Paradox

RICHARD A. HEALEY
AND
GEOFFREY HELLMAN, EDITORS

QUANTUM MEASUREMENT

Beyond Paradox

Richard A. Healey and
Geoffrey Hellman, editors

MINNESOTA STUDIES IN THE PHILOSOPHY OF SCIENCE XVII

MINNESOTA STUDIES IN THE PHILOSOPHY OF SCIENCE

This page intentionally left blank

Minnesota Studies in the PHILOSOPHY OF SCIENCE

C. KENNETH WATERS, GENERAL EDITOR
HERBERT FEIGL, FOUNDING EDITOR

VOLUME XVII

Quantum Measurement: Beyond Paradox

EDITED BY
RICHARD A. HEALEY
AND
GEOFFREY HELLMAN

University of Minnesota Press
Minneapolis
London

Sections of "Varieties of Quantum Measurement," by William G. Unruh, were originally published in *New Techniques and Ideas in Quantum Measurement Theory,* edited by Daniel Greenberger, New York Academy of Science, 1986; reprinted with permission.

Published by the University of Minnesota Press
111 Third Avenue South, Suite 290
Minneapolis, MN 55401-2520
http://www.upress.umn.edu

Printed in the United States of America on acid-free paper

Library of Congress Cataloging-in-Publication Data

Quantum measurement : beyond paradox / edited by Richard A. Healey and Geoffrey Hellman.
p. cm. — (Minnesota studies in the philosophy of science : v. 17)
Includes index.
ISBN 0-8166-3065-8 (hc : alk. paper)
1. Quantum theory—Congresses. 2. Physical measurements—Congresses. I. Healey, Richard. II. Hellman, Geoffrey. III. Series.
QC173.96.Q363 1998
530.12—dc21 98-17341

The University of Minnesota is an equal-opportunity educator and employer.

10 09 08 07 06 05 04 03 02 01 00 99 98 10 9 8 7 6 5 4 3 2 1

Contents

The Bare Theory Has No Clothes

1. What Makes the Bare Theory So Seductive?

Consider quantum theory without the collapse postulate. No experiment has ever disconfirmed its statistical predictions. So one might reasonably bet that no experiment ever will. If not, every possible state of anything in the universe, from the spin states of electrons to the states of consciousness of sentient beings, must evolve in time according to Schrödinger's equation. A theoretician's paradise . . . except for that nagging problem of how to make sense of what the theory says will happen when we try to confirm its statistical predictions.

The problem is worth rehearsing. Suppose m is a device that reliably records whether an electron e's spin is *up* or *down* along some specified direction (for simplicity: without disturbing that spin). The unitary evolution characterizing the $m + e$ interaction will therefore map:

$$\begin{aligned} |up\rangle_e \, |\text{`?'}\rangle_m &\rightarrow |up\rangle_e \, |\text{`}up\text{'}\rangle_m, \\ |down\rangle_e \, |\text{`?'}\rangle_m &\rightarrow |down\rangle_e \, |\text{`}down\text{'}\rangle_m, \end{aligned} \tag{1}$$

where $|\text{`?'}\rangle_m$ is the ready-to-measure state of $m$, and $|\text{`}up\text{'}\rangle_m$ and $|\text{`}down\text{'}\rangle_m$ are two orthogonal recording states of m that distinguish whether e's spin is *up* or *down* in the direction of measurement. (The quotation marks signify that, for present purposes, we can remain agnostic about precisely what physical quantity of m these three states are eigenstates of.) If m is set up to do its job when e is in some arbitrary superposition of *up* and *down* spin states:

$$c_1 \, |up\rangle_e + c_2 \, |down\rangle_e, \tag{2}$$

the above evolutions and the linearity of the Schrödinger equation entail that the final state of $m + e$ will be:

$$c_1 \, |up\rangle_e \, |\text{`}up\text{'}\rangle_m + c_2 \, |down\rangle_e \, |\text{`}down\text{'}\rangle_m. \tag{3}$$

Quantum theory dictates that, on looking to see what recording property m has, there is a chance of $|c_1|^2$ that an experimenter, let's call her Eve, will see '*up*,' and a chance of $|c_2|^2$ that she will see '*down*.'

Now if the Schrödinger equation is universally valid and Eve is a competent observer, we should be able to model the acquisition of her belief about the spin state of e (which she acquires through her belief about m's recording property) by a unitary mapping of the same general form as (1). Then Eve's looking at m will generate the state:

$$\begin{aligned} &c_1\ |up\rangle_e\ |\text{`}up\text{'}\rangle_m\ |\text{`Believes } e\text{-spin } up\text{'}\rangle_{\text{Eve}}\ + \\ &c_2\ |down\rangle_e\ |\text{`}down\text{'}\rangle_m\ |\text{`Believes } e\text{-spin } down\text{'}\rangle_{\text{Eve}}, \end{aligned} \tag{4}$$

where the two states of Eve in this superposition[1] are vectors lying in distinct eigenspaces of whatever physical quantity of Eve's brain it is that records information gathered through her senses (again, the quotation marks allow us to remain agnostic about what physical quantity that is).

At this point, for the statistical predictions of quantum mechanics to have something to refer to, it seems that one must say that, although (4) is the full quantum state of e + m + Eve, in any particular case when Eve looks at m, just one of (4)'s terms represents what Eve will actually come to believe about e's spin. Indeed, statistical predictions aside, something like this *must* be said if we are to reconcile the theory with the fact that experimenters like Eve always take themselves to have definite beliefs about what their measurement devices indicate.

But the trouble is that the standard way of thinking about superpositions in quantum mechanics prohibits attributing a definite value to any observable of a system whose eigenstates are superposed by the system's quantum state, as Eve's belief states are in (4). To say anything more than that by way of underwriting the definiteness of Eve's e-spin belief with hidden variables, worlds, minds, or what-not is standardly taken to be adding to the theory, and that's standardly taken to be a bad thing to do. So it appears that standard thinkers—at least those who want to uphold the universal validity of the Schrödinger equation without introducing the collapse postulate—have backed themselves into a corner on the issue of whether experimenters like Eve have definite beliefs about their measurement outcomes (not to mention the statistics of those outcomes).

But maybe not. David Albert (1992, 116–19) has suggested that if the issue is simply whether it is possible to recover our everyday sense that experimenters hold definite beliefs about the things with which they interact, then that *can* be done with standard thinking using what Albert calls the Bare theory.

The Bare theory promises to be just what its name suggests: a theory on which the Schrödinger equation *is* universally valid, the standard way

of thinking about superpositions *is* correct, and that is all there is to say. So there are no collapses, variables, worlds, or minds postulated to save the definiteness of Eve's belief in (4).[2] The Bare theory's explanation of why, despite all that, everyone (Eve included) will always *come to believe* that her belief about *e*'s spin is definite, is described by Albert as "amazingly cool."[3]

The explanation goes like this (with minor embellishment). Suppose Eve's partner Adam is determined to figure out whether Eve has a definite belief about *e*'s spin while in state (4). Adam cannot just ask Eve what belief about *e*-spin she has while in state (4), because the linearity of the Schrödinger equation will put her into a superposition of responding to Adam in two different ways, and according to Albert (1992, 117), "it won't be any easier to interpret a 'response' like that than it was to interpret the superposition of brain states in [(4)] that that response was intended to be a *description* of!"

Without automatically inducing a superposition of responses, however, Adam can simply ask Eve: "Don't tell me whether you believe the electron to be [*up*] or you believe it to be [*down*], but tell me merely whether or not *one* of those two is the case; tell me (in other words) merely whether or not you *have* any particular definite belief (not uncertain and not confused and not vague and not superposed) about the value of the [spin] of this electron" (1992, 118). The way to model Eve answering that question within quantum mechanics is, presumably, as follows.

If Eve is an honest and competent reporter of her mental states, then when asked whether she has a definite *e*-spin belief in a state like $|\text{'Believes } e\text{-spin } up'\rangle_{\text{Eve}}$, she should report Yes. For in that case, the Bare theory predicts that Eve's belief *will* be definite, and we can assume that there is nothing else stopping her from honestly and competently reporting that fact. But then the same goes for how Eve will respond if she were asked the same question by Adam while in the state $|\text{'Believes } e\text{-spin } down'\rangle_{\text{Eve}}$; namely, she will again respond with a Yes.[4]

It follows that the correct way to model Eve's response to Adam's question about the definiteness of her *e*-spin belief is in terms of a unitary evolution that maps:

$$
\begin{aligned}
&|\text{'Believes } e\text{-spin } up'\rangle_{\text{Eve}}\,|\text{'Ready to Answer'}\rangle_{\text{Eve}} \\
\rightarrow\ &|\text{'Believes } e\text{-spin } up'\rangle_{\text{Eve}}\,|\text{'Yes, I have a definite } e\text{-spin belief'}\rangle_{\text{Eve}}, \\
&|\text{'Believes } e\text{-spin } down'\rangle_{\text{Eve}}\,|\text{'Ready to Answer'}\rangle_{\text{Eve}} \qquad (5) \\
\rightarrow\ &|\text{'Believes } e\text{-spin } down'\rangle_{\text{Eve}}\,|\text{'Yes, I have a definite } e\text{-spin belief'}\rangle_{\text{Eve}},
\end{aligned}
$$

where the two states of Eve we have introduced denote the physical correlates in her brain of her getting ready to answer, and responding Yes to the question.[5]

But now the cool thing is that if (5) models Eve's response to Adam, the linearity of the Schrödinger equation demands that when she responds to his question in superposition (4), the final state of e + m + Eve will necessarily be:

$$c_1 |up\rangle_e |\text{'}up\text{'}\rangle_m |\text{'Believes } e\text{-spin } up\text{'}\rangle_{\text{Eve}} |\text{'Yes, I have a definite } e\text{-spin belief'}\rangle_{\text{Eve}} + \quad (6)$$
$$c_2 |down\rangle_e |\text{'}down\text{'}\rangle_m |\text{'Believes } e\text{-spin } down\text{'}\rangle_{\text{Eve}} |\text{'Yes, I have a definite } e\text{-spin belief'}\rangle_{\text{Eve}},$$

which is *not* a superposition of Eve giving different answers, but an eigenstate of Eve responding Yes! And so the Bare theory predicts that Yes will definitely be Eve's answer, even when in fact she has no definite state of belief about e-spin according to that theory (that is, according to standard thinking about what it means to be in an uncollapsed superposition of belief states). Eve is "apparently going to be radically deceived even about what *her own occurrent mental state* is" (1992, 118). Albert calls a situation like this a situation in which Eve "effectively knows" what e's spin is (1992, 120).

So the cool thing is that we can use the sparse resources of the Bare theory to show why no two people are ever going to believe anything out of the ordinary about the definiteness of each other's beliefs. Indeed, since it could just as well have been Eve herself who introspects about her own beliefs and inquires as to whether they are definite about e-spin, the Bare theory predicts that experimenters will never believe anything out of the ordinary about their own beliefs either.

It is rather like the situation before the fall: Both Adam and Eve were naked, but to keep the paradise, God ensured that they knew not that they were naked. In this case, the Schrödinger equation is what keeps the paradise:

> That is: maybe (even if the standard way of thinking about what it means to be in a superposition is the right way of thinking about what it means to be in a superposition) the linear dynamical laws are nonetheless the complete laws of the evolution of the *entire* world, and maybe all the appearances to the contrary (like the appearance that experiments have outcomes, and the appearance that the world doesn't evolve deterministically) turn out to be just the sorts of *delusions* which *those laws themselves* can be shown to *bring on*! (1992, 123)

Of course, paradise in the garden of Eden did not last long. And trouble is lurking in the Bare theoretician's paradise too. In fact, this story has at least two problems. We shall draw out the first problem in the next section, and identify the second, the decisive problem, in the section after that.

Albert gives further Bare theory stories to explain why experimenters take immediately repeated measurements to yield the same determinate result, why different experimenters measuring the same observable take

themselves to agree on the outcome, and why experimenters take the measurement result statistics they gather in their laboratories to confirm the statistical predictions of quantum mechanics (1992, 119–23). We believe that those stories have exactly the same problems, and shall indicate why toward the end of this chapter.[6]

2. The Bare Necessities

The first thing to sort out is how much of this story about Eve always taking her beliefs to be definite and responding accordingly is a necessary consequence of the Bare theory. Albert homes in on the correct way to model Eve's response to Adam's question by saying that she, being honest and competent enough to assess her own mental state, would answer Yes if she were, according to the Bare theory, in any quantum state that corresponds to her having a definite *e*-spin belief. Albert then employs that model, defined by equations (5), in a state where the Bare theory blocks Eve from having a definite belief. Finally, he concludes on the basis of the linearity of the Schrödinger equation that, nevertheless, she would still respond Yes. But then it would seem she cannot have been honest and competent after all! Perhaps the correct conclusion is that Albert has stacked the deck in his own favor by not modeling Eve so that she can respond as competently and reliably as possible, whatever her circumstances.

But that conclusion is too quick. It fails to recognize that the Bare theory itself sets definite limits on modeling Eve's response. We summarize these limits in a modest theorem.

Bare Theorem

The following are mutually inconsistent:

1. There is a unitary interaction that models Eve answering Adam's question about the definiteness of her *e*-spin belief.

2. Eve's answers are *always* given honestly and competently on the basis of what her beliefs are—or are not—as determined by the Bare theory's standard thinking about superpositions. So she will answer Yes if in an eigenstate of definite *e*-spin belief, and No if not.

3. The model must allow Adam (or Eve, if she is just introspecting about the structure of her own beliefs) to distinguish situations in which Eve chooses to respond (or conclude) Yes from situations in which she chooses to respond (or conclude) No.

The inconsistency follows easily. Let $|\text{'Huh?'}\rangle_{\text{Eve}}$ be any nontrivial superposition of $|\text{'Believes } e\text{-spin } \textit{up}\text{'}\rangle_{\text{Eve}}$ and $|\text{'Believes } e\text{-spin } \textit{down}\text{'}\rangle_{\text{Eve}}$. Assumptions 2 and 3 entail that the unitary evolution that 1 assumes to exist must, in particular, map:

$$\begin{aligned} &|\text{'Believes } e\text{-spin } \textit{up}\text{'}\rangle_{\text{Eve}}\,|\text{'Ready to Answer'}\rangle_{\text{Eve}} \\ \rightarrow\ &|\text{'Believes } e\text{-spin } \textit{up}\text{'}\rangle_{\text{Eve}}\,|\text{'Yes, I have a definite } e\text{-spin belief'}\rangle_{\text{Eve}}, \\ &|\text{'Huh?'}\rangle_{\text{Eve}}\,|\text{'Ready to Answer'}\rangle_{\text{Eve}} \qquad (7) \\ \rightarrow\ &|\text{'Huh?'}\rangle_{\text{Eve}}\,|\text{'No, I don't have a definite } e\text{-spin belief'}\rangle_{\text{Eve}}, \end{aligned}$$

which it cannot, since that would be to map initially nonorthogonal vectors onto orthogonal ones. (The contradiction continues to hold without the idealization that in responding Eve fails to disturb her own state of belief.)

So from the point of view of the Bare theory, if we understand the theory as requiring 1, 2, and 3 above for a "good" measurement of the definiteness of Eve's *e*-spin belief, there *is no* good way to model Eve's response quantum-mechanically.

Nevertheless, since quantum mechanics is everything on the Bare theory, and we think we *do* ascertain the definiteness of each other's (and, indeed, our own) beliefs via interactions of *some* sort, those interactions must be capable of being modeled by *some* kind of unitary evolution. So 1 is not open to question. And 3 cannot be given up without undermining the possibility of Eve ascertaining and reporting on what the structure of her *e*-spin beliefs is, which is presumably something we *do* think she can do. So 2 needs to be loosened up in some way.

But in what way? Albert takes it that the way to model Eve's response unitarily is via an interaction of type (5). He appeals to the competency Eve would have to report her mental state when it is an eigenstate of some definite belief about *e*-spin. He then rigs the unitary interaction of (5) so that it correlates the answer Yes to those eigenstates,[7] concluding (by linearity) that Eve must be completely *in*competent in reporting her mental state when in a superposition of them.

However, the following parallel line of argument is equally compelling, and equally natural in the context of the Bare theory. Suppose, instead, that Eve is competent to report her mental state when it is *not* an eigenstate of some definite belief about *e*-spin. (After all, why assume a priori that she would somehow be speechless about her *lack* of definite belief?) Let $|\text{'Huh?'}\rangle_{\text{Eve}}$ and $|\text{'Wha?'}\rangle_{\text{Eve}}$, not eigenstates of definite belief about *e*-spin, span the same plane in the Hilbert space that represents Eve's brain states as the eigenstates $|\text{'Believes } e\text{-spin } \textit{up}\text{'}\rangle_{\text{Eve}}$ and $|\text{'Believes } e\text{-spin } \textit{down}\text{'}\rangle_{\text{Eve}}$ do. Then in each of those noneigenstates of *e*-spin belief, Eve should respond No to the question, "Do you have a definite *e*-spin belief?" By

linearity, she will also respond No for any brain state in the plane. And that includes the very states in which the Bare theory guarantees that she *will* have a definite belief, namely, $|\text{'Believes } e\text{-spin } up\text{'}\rangle_{\text{Eve}}$ and $|\text{'Believes } e\text{-spin } down\text{'}\rangle_{\text{Eve}}$!

For example, if

$$\begin{aligned} |\text{'Huh?'}\rangle_{\text{Eve}} &= \frac{1}{\sqrt{2}}|\text{'Believes } e\text{-spin } up\text{'}\rangle_{\text{Eve}} + \frac{1}{\sqrt{2}}|\text{'Believes } e\text{-spin } down\text{'}\rangle_{\text{Eve}}, \\ |\text{'Wha?'}\rangle_{\text{Eve}} &= \frac{1}{\sqrt{2}}|\text{'Believes } e\text{-spin } up\text{'}\rangle_{\text{Eve}} - \frac{1}{\sqrt{2}}|\text{'Believes } e\text{-spin } down\text{'}\rangle_{\text{Eve}}, \end{aligned} \tag{8}$$

then

$$\begin{aligned} |\text{'Believes } e\text{-spin } up\text{'}\rangle_{\text{Eve}} &= \frac{1}{\sqrt{2}}|\text{'Huh?'}\rangle_{\text{Eve}} + \frac{1}{\sqrt{2}}|\text{'Wha?'}\rangle_{\text{Eve}}, \\ |\text{'Believes } e\text{-spin } down\text{'}\rangle_{\text{Eve}} &= \frac{1}{\sqrt{2}}|\text{'Huh?'}\rangle_{\text{Eve}} - \frac{1}{\sqrt{2}}|\text{'Wha?'}\rangle_{\text{Eve}}. \end{aligned} \tag{9}$$

If Eve (competently) responds No to the question about the definiteness of her e-spin belief in the states $|\text{'Huh?'}\rangle_{\text{Eve}}$ and $|\text{'Wha?'}\rangle_{\text{Eve}}$, by linearity she will (*in*competently) respond No to the question in the states $|\text{'Believes } e\text{-spin } up\text{'}\rangle_{\text{Eve}}$, $|\text{'Believes } e\text{-spin } down\text{'}\rangle_{\text{Eve}}$.

So two completely parallel arguments for how to model Eve's response yield utterly incompatible results and incompatible stories about the way in which Eve is deluded. The Bare theory per se does not tell us which is the correct model, because it contains no principle that stipulates in what ways experimenters should be competent to report their beliefs, and in what ways they should not. In other words, the Bare theory contains no prescription that says that when Eve reflects about her own beliefs, Eve's brain is hard-wired so that Albert's unitary evolution (5) gets turned on inside her brain, as opposed to the unitary evolution just described. The Bare theorist cannot treat evolution (5) as a kind of calibration condition that defines what we mean by Eve's response acting as a good measurement of the structure of her own beliefs. For the evolution which induces her to respond No in both the $|\text{'Huh?'}\rangle_{\text{Eve}}$ and $|\text{'Wha?'}\rangle_{\text{Eve}}$ states is based on an equally good calibration condition, given that under the Bare theory the No response would be the correct one for those states.

However, this is still not enough to reject the Bare theorist's story out of hand, for there is an easy way Albert can grant the point just made. No interpretation of quantum mechanics should ever be called upon to decide a priori how to model a given interaction. So, in this case, we are free to find a model that fits the data that we never find people reporting that they fail to have definite beliefs about things when they are in eigenstates of definite belief—assuming that is an empirically established (or, at any rate, estab-

lishable) fact. Then, the ability of the Bare theory to at least represent (if not strictly deduce) the fact that experimenters never take their beliefs to be indefinite is perhaps all that should be required.

But before this aspect of the Bare theory is shouted from the rooftops, we would do well to put it into perspective. Compare what remains of the Bare theory's explanatory strategy to the strategy adopted by nonstandard interpretations of no-collapse quantum mechanics that inject variables, worlds, or minds into the discussion to make definite Eve's beliefs in a superposition such as (4).

The chief reason why standard thinkers about superpositions dislike the nonstandard strategy is that it requires certain observables—such as the observable whose eigenstates underpin Eve's beliefs—always to have well-defined values, even when the observable's eigenstates are superposed by the quantum state of the system. Standard thinkers do not accept that sort of discrimination between physical observables: as far as the formalism of quantum mechanics is concerned, all of its observables are on a par. To single out whatever observable it is that grounds our beliefs as an observable that it always makes sense to ascribe a definite value to regardless of quantum state is regarded as irredeemably ad hoc. There is just nothing special about that observable as compared to any other; nothing "preferred" about the set of belief eigenspaces in the Hilbert space representing Eve's brain over any other observable's eigenspaces in that space. (And, of course, if we assume all observables have simultaneously definite values all the time, that all sets of eigenspaces are on a par with respect to determinacy, then we run into well-known difficulties with what we mean by saying that every observable has a determinate value [see note 10].)

In that regard, however, the Bare theory's strategy for dealing with the problem of what experimenters take themselves to believe on the basis of their measurement devices is now no less objectionable than the nonstandard approach. The nonstandard preferred-observable approach would say that it just so happens in our world that "belief" is the kind of physical observable for which it always makes sense to speak about having one belief or another (where that may, in purely physical terms, amount to saying that the particles in experimenters' brains take up one configuration over another, if what we call beliefs turn out to be nothing other than certain positions being taken up by certain brain particles). The Bare theorist has to say that it *just so happens* in our world that the brain is hard-wired to give us the kind of beliefs we take ourselves to have about our beliefs.

In fact, the analogy between the two explanatory strategies is far closer than that. The difference between the unitary evolution that would make Eve respond Yes and the one that would make her respond No is just a difference in choice of basis: either we decide that she is in fact competent to

report the definiteness of her belief in both of two orthogonal belief eigenstates, or we decide that she is in fact competent to report her lack of definite belief in both of two orthogonal *non*eigenstates of belief that span the same plane. The comparable choice on a preferred-observable approach is between Eve's belief observable as determinate versus, say, the observable with $|\text{'Huh?'}\rangle_{\text{Eve}}$ and $|\text{'Wha?'}\rangle_{\text{Eve}}$ as eigenstates. The formal structure in quantum mechanics that needs to be distinguished by the Bare theory to get a fix on modeling observers' reports *is exactly the same structure* preferred-observable approaches exploit, namely a preferred basis in the subspace spanned by $|\text{'Believes } e\text{-spin } up\text{'}\rangle_{\text{Eve}}$ and $|\text{'Believes } e\text{-spin } down\text{'}\rangle_{\text{Eve}}$ picked out by those two vectors.

So despite its name, the Bare theory cannot get by with less than nonstandard preferred-observable approaches do in explaining why experimenters report themselves to have definite beliefs. It cannot be claimed that the Bare theory's choice is determined by empirical facts any more than that can be claimed for the preferred-observable approaches. The fact that experimenters respond Yes in belief eigenstates, and also respond like that in any state that is a linear superposition of those eigenstates, is exactly the sort of empirical fact preferred-observable approaches can point to in justifying their own particular preference, namely, for the belief observable as the one that's determinate.

3. The Bare Theory Exposed

So far we have argued that if the Bare theory's story about how experimenters come to believe/report that their beliefs are definite is granted, then it is no better off than preferred-observable approaches with respect to the structure in the quantum formalism that needs to be distinguished. We now turn to arguing that the Bare theory is actually worse off because it does not succeed in explaining all that needs to be explained. The final section of the chapter will point out how the preferred-observable approach succeeds exactly where the Bare theory falls short. The upshot will be that this entire exercise concerning the Bare theory supplies a nice (even if somewhat scholastic) argument for a preferred-observable approach to universal no-collapse quantum mechanics.

It will be important, first, to be absolutely clear about what the standard thinking about quantum states is, since our contention will be that Albert does not follow that thinking to its logical conclusion.

There is no difficulty in capturing what standard thinking is about eigenstates. If a physical system is in an eigenstate of an observable, then that observable possesses the corresponding eigenvalue (or at least the system

has the surefire disposition to behave as if that value were a property it possesses). In superpositions of two (or more) distinct-eigenvalue eigenstates of an observable, things are more delicate.

Suppose the eigenstates' eigenvalues are x and y. Then standard thinking entails that it would not be right to say in the superposition that the observable has value x, nor would it be right to say it has value y, nor would it be right to say it has *both* values, nor would it be right to say it has *neither* value. This is Albert's own understanding of standard thinking about superpositions (1992, 11, 79n). However mind-boggling these assertions might appear, they are the rules of the game that need to be followed if the Bare theory is to be entertained.

Now in our thought experiment about Eve, Albert makes a conceptual distinction between her brain as a repository of beliefs about *e*-spin, and the aspects of her that reflect on her beliefs about *e*-spin and then report on her reflections if she is asked questions about her beliefs. The fact that she verbalizes her beliefs about her *e*-spin beliefs is unimportant, although at one point Albert suggests it might be important: "[Eve] is necessarily going to be convinced (or at any rate she is necessarily going to *report*) that she *does* have a definite particular belief" (1992, 118). It cannot possibly be important, because if she does not *think* herself to have a definite *e*-spin belief, then that is presumably going to be a problem for the Bare theory regardless of what she reports. Furthermore, the very act of her reporting that she has a definite *e*-spin belief should presumably induce her to give assent to the content of her report, if the way we are modeling Eve as a quantum-mechanical automaton is going to have any relation to our own experience that we (in normal circumstances) don't bear witness to things we aren't prepared to accept ourselves.

So we can view the eigenstates of Eve's *reporting* Yes that we introduced earlier as eigenstates of her *reflecting* Yes (but not necessarily verbalizing the answer) without compromising the Bare theory in any way. Then what Albert needs is a conceptual distinction between Eve's brain as a retainer of beliefs about things such as *e*-spin and Eve's brain as it functions in reflecting about the structure of those beliefs. For easy reference, call these two faces of Eve the "inner" and "outer" Eve respectively.[8] And let's grant to Albert that they will interact in exactly the way he postulates, that is, in accord with equations (5).

For Albert to make his point that the Bare theory predicts that Eve will be deluded about the definiteness of her *e*-spin beliefs, he needs to establish that the outer Eve draws a conclusion at variance with the Bare facts of the matter about the inner Eve. And for both inner and outer Eve, Albert must employ standard thinking to determine the Bare facts of the matter about Eve's inner beliefs, and about her outer beliefs about her inner beliefs.

The inner facts of the matter in the belief superposition are, in particular, that Eve doesn't believe *up* and she doesn't believe *down*. These are quite specific negative claims about her situation. What we are leading up to is that Albert unjustifiably pulls back from drawing a similar kind of negative conclusion when outer Eve has the potential of getting into the same kind of superposition.

When outer Eve reflects upon whether she has exactly one of the two beliefs in the set {Believes *e*-spin *up*, Believes *e*-spin *down*}, she always concludes Yes (on the assumption that a particular unitary interaction, namely (5), is singled out as modeling this act of reflection). So if the Bare theory is true, Albert is right to observe that Eve is deluded when answering Yes in the superposition. But if the Bare theory *is* true, we can also ask what it will predict when Eve attempts to reflect upon what belief about *e*-spin she has. Since she would then get into a superposition of believing that she believes *up* and believing that she believes *down* (assuming she is good at ascertaining her specific beliefs in the eigenstate cases), under the Bare theory she will be unable to specify *which* of the two beliefs she takes herself to hold.

Recall what Albert says about that when Eve was being asked to report to Adam her reflections about what *e*-spin belief she has while in the superposition: "it won't be any easier to interpret a 'response' like that than it was to interpret the superposition of brain states in [(4)] that that response was intended to be a *description* of!" True enough. But the general difficulties of interpreting superpositions notwithstanding, the Bare theory has something quite specific to say about Eve getting into a superposition of believing that she believes *up* and believing that she believes *down*. The Bare theory says that outer Eve will *not* believe that she believes *up* and will *not* believe that she believes *down*, in just the same way that it says that inner Eve does not believe *up* and does not believe *down*.

The point is that we don't merely take ourselves to have "definite beliefs" in everyday life. After performing spin measurements on electrons, we sometimes take ourselves to believe *up* and other times take ourselves to believe *down*, depending on how our experiments go. The Bare theory can only explain *that* when we perform measurements on electrons in spin eigenstates—not in superpositions. The issue that seems to have been lost sight of in all of this is that we *also* commonly believe that we have the *grounds* for believing that we have definite beliefs, in the sense that on any given occasion we are able to reflect on what our belief is!

Let's put it another way. Earlier in his book (1992, 78–79), Albert takes the problem of interpreting a universal quantum mechanics without the collapse postulate to be that the theory predicts we should have no particular

belief in a superposition when, by "direct introspection," we apparently do. His strategy for the Bare theory is to see if we can get it to explain why we take ourselves to have a definite belief in a superposition. But Albert has merely solved a watered-down version of the problem. The problem is not just to explain why we take our beliefs about things like *e*-spin to be definite in some noncommittal sense, but why we take ourselves to believe *the specific things that we do apparently believe* about *e*-spin! In trying to get the Bare theory to explain *that*, it will (as Albert himself admits) land outer Eve into the same sort of devastating superposition that inner Eve is in. But that fact is hardly irrelevant to the Bare theory's inability to explain what we take our specific beliefs to be.

The Bare theory's story about apparently definite beliefs is so seductive (as each of the present authors knows from personal experience) that anticipating responses to this quite elementary criticism we have just made is obligatory.

Note, first, that we are not accusing the Bare theory of any flat-out logical contradiction with respect to Eve's beliefs. She does believe she has exactly one of two *e*-spin beliefs, even while:

she doesn't believe that she has the *up* belief; and

she doesn't believe that she has the *down* belief. (10)

This *is* an extraordinary state of affairs, the likes of which we just don't ever seem to find ourselves in. But surely it is no more logically contradictory than asserting the provability within some formal system of "$G \vee \neg G$"—where G is a Gödel sentence of that formal system—while denying the provability of G and the provability of $\neg G$.

On the other hand, we are *not* making the move from (10) to:

she believes that she *doesn't* have the *up* belief; and

she believes that she *doesn't* have the *down* belief. (11)

The pair of conclusions about Eve in (11) obviously do logically contradict her taking herself to have exactly one of the two *e*-spin beliefs. But to draw either of them on the Bare theory when Eve is in the superposition, outer Eve would have to get into an eigenstate of belief about something particular about *e*-spin that she doesn't believe. There has been no suggestion of that possibility here.[9]

Nor do we need to establish anything like that. It suffices for our critique that the Bare theory is unable to represent the fact that we do sometimes take ourselves to believe a specified thing such as *up* when, according to the theory, we are in a state like (4). That is, it suffices to point out

that the Bare theory's standard thinking leaves Eve with *no impression whatsoever* about whether or not she believes *up* on any particular occasion, even granting that she is mistakenly convinced that her *e*-spin beliefs are definite on all occasions.

Now we have granted Albert that Eve will mistakenly (but without being aware of her mistake) believe that she either believes *e*-spin to be *up* or believes it to be *down* (which is, of course, what believing that her *e*-spin belief is definite is all about). But then if she is rational and knows in advance that she is measuring a spin-$\frac{1}{2}$ particle, won't she also sometimes conclude that she believes *up*, and other times conclude that she believes *down*, even while she is in the superposition? That seems to be the obvious way to read Albert's claim that Eve "effectively knows" *e*'s spin.

But it's not on, for the simple reason that we cannot just endow Eve with magical powers of reasoning that transcend her states of belief as determined by the Bare theory. If we say that on some occasion while in the superposition Eve draws the conclusion that she believes *up* from her sense that her beliefs about *e*-spin are definite, we are saying something that flatly contradicts the Bare theory unless she can get into an eigenstate of drawing that conclusion. But the only way she can do that is if she is actually in an eigenstate of believing *up* and *not* in the superposition, as we have already emphasized.[10]

Another way to avoid our criticism might be for Albert to introduce a third level of Eve (giving us now the *three* faces of Eve) that reflects upon what her reflections about her beliefs are. When asked in the superposition if her reflection on what her *e*-spin belief is agrees with her actual *e*-spin belief, Eve can be modeled so as to respond (you guessed it) *Yes*. In fact, this is exactly what Albert uses to explain the repeatability of measurements, and the agreement between different experimenters who measure the same electron as to what its spin is (1992, 119–20).

But that's not on either. Granted the Bare theory can explain intersubjective agreement between experimenters about what their measurement outcomes are. But what cannot be explained is why sometimes the experimenters will take that agreement to be established because both their results were *up*. Mere "agreement" is not enough. The same goes for the three faces of Eve. Even if her third face is convinced that her second- and first-level selves are in accord about her particular *e*-spin belief, that does nothing to give her any sense of what those selves have agreed to. Indeed, agreement tests between first-, second-, third- . . . nth-level Eve could never be used as the factual basis for her feeling or belief that she believes *up* after some particular measurement trial, simply because in superposition (4) she can never get into any sort of eigenstate that makes reference to just the one particular belief *up*.[11]

At this point the reader might have become suspicious about our introduction of inner and outer Eve, and that our troubles with the Bare theory stem from that naive conception of Eve. All there is to Eve is her beliefs and the way she reports on them, and nothing in between! That is certainly more faithful to Albert's *text*. But it does little to mute the problem; it only changes the terms in which the problem needs to be stated. The problem would then become that, even if Eve operationally qualifies as a person convinced of the definiteness of her *e*-spin belief, in the superposition she will never be able to qualify as a person convinced she has the particular belief that *e*-spin is *up*. Talking about an introspective side of Eve versus what an external observer would infer about what she takes her mental state to be changes nothing. And this is so even if we demand—as Albert's phrasing of Adam's question does—that Eve give verbal assent to the fact that her *e*-spin belief is "not uncertain and not confused and not vague and not superposed."

So much for anticipating possible responses to our main criticism. The same criticism undermines the story Albert runs to explain why, on the Bare theory, experimenters take themselves to have confirmed the statistical predictions of quantum mechanics through their measurements (1992, 120–23). In the limit of an infinite sequence of measurements, an experimenter's brain will get into a quantum state that lies in an eigenspace of the relative frequency operator, with an eigenvalue equal to the quantum-mechanically predicted frequency for the measured observable. And so, according to the Bare theory, experimenters will take themselves to have obtained sequences that agree, in their long-run frequencies, with quantum statistical predictions.

But experimenters as we know them seem to believe far more than that. They take themselves to have obtained particular sequences with particular characteristics. Merely getting into an eigenspace of the relative frequency operator will not ensure *that* on the Bare theory, since there are numerous distinct sequences that can yield the same long-run frequencies. Albert's own problem with this story is that it is too idealized. Determinateness with respect to the relative frequency operator is only achieved in the limit of an infinite number of measurements (1992, 124–25). But even the Bare theory's idealized experimenter is unlike any we are ever likely to meet (or be).

A similar point is made by Stein (1984) against Geroch's (1984) reading of Everett's interpretation (which is essentially the Bare theory; see notes 2 and 3). Stein asks how we ordinarily go about checking that measurement outcomes conform to the relative frequencies predicted by quantum mechanics, and answers (1984, 641, emphasis in the original): "We do so by performing the experiments and *noting and counting their outcomes*." We

too have argued that *that* part of our ordinary experience cannot be accounted for on the Bare theory. However, Stein goes further than he should (or needs to) in suggesting that Geroch's observers are committed to the exclusive use of physically special apparatus designed to record only statistical information and "obliviate" particular outcomes:

> Geroch would have us (conceptually at least) interpose, between the experimenters and the *real* observation, apparatus that is designed just to detect whether or not the predicted frequencies (within the predicted range of variability) have occurred: what we may call "obliviating apparatus," designed expressly to ignore everything about the series of experiments that was not predicted or precluded with practical certainty by quantum mechanics. The idea that physicists might actually proceed in this way—spend money and time on instrumentation (automatic recording apparatus to feed data to computers, programs designed to report only the predicted aspects of the data and to destroy all records of details that were not predicted)—simply in the interest of conformity to the Everett-Geroch interpretation, seems worthy of Swift's Academy of Lagado. (1984, 642)

In spite of the initial qualification, it appears that Stein regards the obliviating apparatuses required by the Bare theory (in Albert's or Geroch's version) as more than merely conceptual. This goes too far, since the Bare theory has no need to postulate devious devices designed specifically to lose or erase information about outcomes while retaining statistical information. The fact that an apparatus, or an observer's mind, only retains long-run statistical information follows quite straightforwardly from the observation that information is registered only if the state of the registering system is an eigenstate of registering it! This is a direct consequence of the standard thinking about superpositions employed in the Bare theory and needs no further special physical assumptions about the design of measurements, beyond those commonly granted in standard analyses of quantum measurement. So while we agree with Stein's starting point, we believe he overstates the case.

Note, finally, that an idealized experimenter measuring an ensemble of electrons in superposition (2) will, in the limit, approach an eigenstate of answering Yes to the question, "Did you sometimes take yourself to believe *up* on the basis of your measurement result?" So isn't it the case, after all, that the Bare theory *can explain* why we sometimes take ourselves to believe *up* (and other times *down*) on the basis of our measurements? No. The alleged explanation here trades on an ambiguity in the word "sometimes." What's important is not just whether at some point or other (sometimes) in the experimenter's measurement outcome sequence

an *up* was taken to be registered, but whether on (say) trial number 13 (some particular time) the experimenter took him- or herself to have measured *up*. Again, getting into an eigenspace of the relative frequency operator will not suffice for a definite answer to *that* question.

We conclude that the Bare theory fails to explain why we take ourselves to have specific beliefs about measurement outcomes or outcome sequences in generic experimental situations. The only conclusion left is that if the Bare theory is true, then in states like (4) we really *never* believe that we believe *up* nor do we ever believe that we believe *down* after (any single!) measurement of an electron's spin. The same goes for believing that we believe anything particular about measurement outcome sequences beyond what their long-run frequencies are (and perhaps other features of outcome sequences definable in the long run, such as their randomness). If that's "What It Feels Like to Be in a Superposition" (Albert 1992, 112), then the Bare theory bears far less on the actual experiences of actual experimenters than we have been led by Albert to believe.

4. A Fig Leaf for the Bare Theory

Fortunately, explanatory adequacy with regard to what we take ourselves to believe is quickly restored on a preferred-observable approach to universal, no-collapse quantum mechanics that simply denies standard thinking about superpositions from the very beginning.[12]

On such an approach, all the endless problems about beliefs about beliefs and responses about beliefs engendered above are cut off at their roots. Because any observable representing beliefs is granted the status of an observable that is always determinate, in superposition (4) inner Eve will believe $|c_1|^2$ of the time that *e*-spin is *up* and $|c_2|^2$ of the time that *e*-spin is *down*. Outer Eve will believe that she has definite beliefs, in agreement with the Bare theory. But she will also take herself to believe *up* $|c_1|^2$ of the time and to believe *down* $|c_2|^2$ of the time. And this is exactly in accord with *all* the things experimenters take themselves to believe. They will not be left (unknowingly) deluded, but more important, not be left without a sense of what they take their specific beliefs about *e*-spin to be.

So making belief the preferred observable (whatever the physical basis of our beliefs turns out to be) is exactly the fig leaf the Bare theory needs to fill the gap between guaranteeing that experimenters have the sense that their beliefs are definite, and giving them the specific feelings they get about the specific beliefs they hold.

Albert's Bare theory argument concerning what Eve takes herself to believe, the argument that we have spent the bulk of this paper analyzing,

goes back at least as far as the paper by Albert and Loewer (1988). Interestingly, they do not treat the Bare theory as a possible end in itself, but use it as a springboard to draw a conclusion that is the same in spirit as our own.

Albert and Loewer use the Eve scenario to argue that if we maintain the idea that experimenters never report that their beliefs are definite when that is in fact false, then their belief states cannot supervene on a superposition like (4), but must be specified in addition to it (1988, 205). In other words, accurate reporting requires a preferred-observable approach with the observable that underpins Eve's beliefs always taken to be determinate independently of the quantum state of Eve.[13]

We agree with Albert and Loewer's conclusion that the Eve scenario most naturally points to a preferred-observable approach, and that assuming reports of belief definiteness are accurate (and there is no collapse) entails that conclusion. But we have shown that one can arrive at that conclusion without any such assumption. Our point has been that the Bare theory is too bare to account for the specificity of what we take ourselves to believe, and that *actual* determinateness of beliefs (as opposed to *effective* determinateness) is necessary to account for that specificity.

The Bare theory's Adam and Eve fall short of what we take ourselves to be like. Having eaten of the fruit from the tree of the knowledge of *up* and *down*, we know what's up.

Notes

We would like to thank John Bell, Jeff Barrett, Geoffrey Hellman, Tim Kenyon, and Steve Weinstein for helpful discussions (without committing them to endorsing anything that has been argued here). R. C. acknowledges support from the Vice President of Research at the University of Western Ontario, and B. M. acknowledges a National Science Foundation Graduate Research Fellowship.

1. Writing down a ket like $|$'Believes *e*-spin *up*'$\rangle_{\text{Eve}}$ carries with it no commitment to there being a unique way to instantiate a state of Eve physically that would count (by operationalist or other criteria) as a state of her believing *up*. Furthermore, the quantum-mechanical observable that underpins "belief," whatever it is, is bound to be degenerate simply on the grounds that any particular belief, like the belief that *e*-spin is *up*, may be held concurrently with a host of other beliefs about other topics—beliefs that will vary from one experimenter to another.

2. Albert (1992, 124) sees the Bare theory as an attractive way to elaborate Everett's (1957) ideas as a "one world" theory that preserves the standard way of thinking about superpositions (that is, a theory based on the "eigenvalue-eigenstate link" as the criterion for value-definiteness or determinateness). While Albert cites Lockwood (1989), who follows Deutsch (1985), Geroch (1984) has also independently proposed an interpretation of Everett that is very much along the same lines. For more recent technical and conceptual

elaborations of this sort of "bare" interpretation of Everett, see Barrett (1994, 1995, 1996). Note that these interpretations differ from Bell's (1987) "one world" interpretation of Everett. Bell's interpretation—as he puts it (1987, 133), "the pilot-wave theory without the trajectories"—is a preferred-observable interpretation of the type considered in section 2, with position in configuration space as the preferred, always determinate observable. We shall not be concerned here with the murky topic of Everett exegesis.

3. Geroch (cf. previous note) puts the aim of this explanatory project well: "what must be accounted for, if anything, is, not the specific classical outcomes deemed to have occurred for a specific experiment, but rather the general human impression that classical outcomes do occur" (1984, 628–29). But, unlike Albert, Geroch does no more than sketch how one might try to complete the project, and remains agnostic about its feasibility. Thus he continues: "This problem may well be soluble, but is probably beyond our present abilities; and, in any case, is basically not a problem in quantum mechanics" (1984, 629).

4. It is not important for what follows that the particulars of *how* Eve responds be the same in the two eigenstates of her *e*-spin belief (that is, responding with a whisper or a scream, via one set of neurons firing in her brain over another set that could equally well do the job, and so on). It is only important that the *content* of her response be the same. In other words, if we consider the subspace in the Hilbert space of Eve spanned by all those states of her that we are willing to count (by operational or other criteria) as states in which the content of her response is Yes, then what's important is that the two occurrences of the ket $|$'Yes, I have a definite *e*-spin belief'$\rangle_{\text{Eve}}$ in (5) both be vectors that lie in that subspace. They need not generate the same ray, even though we have not taken the trouble to reflect that possibility in our notation.

5. For present purposes, we can think of Eve's physical state as an element of the tensor product of the Hilbert space that represents the features of her brain that register beliefs about *e*-spin with the Hilbert space associated with the parts of Eve involved in reflecting on what her beliefs are and responding to questions about them. It does not matter whether the part(s) of Eve's brain associated with the relevant beliefs are individuated neuroanatomically or functionally, that is, whether a Token or Type physicalism is taken to apply. And our case against the Bare theory is unaffected even if one eschews mind-brain identity in favor of the neurological causes of mere belief *reports*.

6. It should perhaps be pointed out that Albert is not entirely uncritical of these stories himself (1992, 124–25). But once our "decisive" problem is laid out, it will become clear that we believe Albert should have been critical of these stories at a far earlier stage in his exposition.

7. It might have occurred to the reader that interaction (5) physically amounts to nothing more than a measurement of the value of the projection operator onto the subspace spanned by $|$'Believes *e*-spin *up*'$\rangle_{\text{Eve}}$ and $|$'Believes *e*-spin *down*'$\rangle_{\text{Eve}}$ and, as such, carries little information about Eve's physical state beyond the fact that it lies in that subspace. But Albert is well aware of this (see his particle in a box analogy [1992, 117n]). Albert's point is that Eve in the superposition will be responding to Adam's question in exactly the same way as she would if she were in an eigenstate of *e*-spin belief (at least with regard to the content of her answer; see note 4), and therefore she (and others) will be given all the same impressions about her state of mind in the two cases. What purely physical meaning one assigns to all of this, in terms of what observable of Eve has actually been measured via (5), would presumably be regarded by Albert as beside the point.

8. This really is for ease of reference: we are not suggesting that Eve has multiple personalities, but only that she is capable of forming beliefs and also (second-order) beliefs about her beliefs. That much is necessary to get the story about Eve's delusion about her beliefs off the ground.

9. Evidently, the closest one can get to (11) on the Bare theory is to have Adam ask Eve the question, "Do you *lack* exactly one of the two beliefs in the set {Believes *up*, Believes *down*}?" Eve will (presumably) respond Yes to this in the superposition if that was her response to the original question about *having* exactly one of those two beliefs (that is, setting aside the previous section's worries about Yes versus No). But clearly answering Yes to both questions is logically consistent. Nevertheless, the Bare theory again remains explanatorily inadequate with respect to the specific things about *e*-spin we take ourselves *not* to believe subsequent to a particular measurement trial on an electron's spin.

10. The fact that superposed Eve is never able to take herself to believe *up* or take herself to believe *down*, even though she takes herself to believe exactly one of those, is analogous to a well-known difficulty faced by realistic quantum logic (as defended, for example, by Putnam [1975]) in its attempt to render meaningful statements such as "Every observable has a definite value." Recall that the difficulty is that, even though the proposition "Observable O has a definite value" may be identified with the disjunction "o_1 or o_2 or . . . or o_n"—where (for any i) o_i is the proposition "Observable O has value o_i"—the truth of this disjunction, which is a tautology in quantum logic, turns out to be compatible with the falsity of all of its disjuncts. What makes this a difficulty for realistic quantum logic is that it waters down the very realism of definite values that logic was designed to prop up. By analogy, what we have argued against the Bare theory is that Eve's impression of the definiteness of her *e*-spin belief, expressed through her belief in a disjunction over her believing *up* versus *down*, necessarily yields only a watered-down representation of our actual experience, because Eve's belief cannot be accompanied, on any given occasion, by belief that a particular one of those disjuncts is true. However, there is a downside to this quantum logic analogy: we do not (and need not) claim that the Bare theory forces Eve (or Adam) to understand her answer to Adam as a conscious endorsement of quantum over classical logic. Indeed, just as Adam can ask Eve what she thinks about the definiteness of her *e*-spin belief and will always elicit a Yes response, he will always elicit that same response by asking her whether she takes the structure of her beliefs to conform to the logic of the classical "or"!

11. Jeff Barrett (in correspondence with the authors) has suggested another possible way to bring the Bare theory closer to endowing Eve with the (illusory) sense that her *e*-spin belief has the content that it should have. Let Eve make spin measurements on three identical electrons: the first in an *up* eigenstate, the second in a *down* eigenstate, and the third in a superposition thereof. Again, Eve will believe she formed a definite belief about all three measurement results. But, furthermore, she will believe that the result of her last measurement was indistinguishable (in its content and the impression it left on her) from the result of exactly one of her first two measurements! Alas, this new "agreement argument" is still too weak to explain why it is that we think we can specify which of the first two measurement results the third agreed with. Put another way, the Bare theory cannot vindicate the belief that the third result was indistinguishable from the first result, or the belief that it was indistinguishable from the second. Beliefs of this sort may well be mistakenly acquired in our world, but the Bare theory is still at a loss to explain why.

12. Note that preferred-observable interpretations, as we understand them, include the modal interpretations discussed by several contributors to this volume. In a modal interpretation formulated along the lines of the versions originally proposed by Kochen (1985) and by Dieks (1989), the preferred observable is derived from the quantum state and is therefore time-dependent, as opposed to a preferred-observable interpretation with a fixed, always determinate, observable. For a detailed analysis comprehending both these ways of choosing a preferred observable, see Bub and Clifton (1996). The situation is more complicated in Healey's (1989) modal interpretation; see Clifton (1996) for further discus-

sion. Bacciagaluppi, Elby, and Hemmo (1996) have shown how the determinateness of observables representing beliefs in the Kochen-Dieks modal interpretation requires decoherence arguments.

13. Albert and Loewer go on to argue that the best preferred-observable approach involves invoking "many-minds," but we do not wish to enter into that discussion here. Suffice it to say that we believe everything presently under discussion is neutral with respect to how one makes the preferred observable preferred, whether it be in terms of (one or many) hidden variables, worlds, or minds.

References

Albert, D. 1992. *Quantum Mechanics and Experience*. Cambridge, Mass.: Harvard University Press.

Albert, D., and B. Loewer. 1988. "Interpreting the Many Worlds Interpretation." *Synthese* 77: 195–213.

Bacciagaluppi, G., A. Elby, and M. Hemmo. 1996. "Observers' Beliefs in the Modal Interpretation of Quantum Theory," manuscript.

Barrett, J. 1994. "The Suggestive Properties of Quantum Mechanics without the Collapse Postulate." *Erkenntnis* 41: 233–252.

———. 1995. "The Single-Mind and Many-Minds Versions of Quantum Mechanics." *Erkenntnis* 42: 89–105.

———. 1996. "Empirical Adequacy and the Availability of Reliable Records in Quantum Mechanics." *Philosophy of Science* 63: 49–64.

Bell, J. S. 1987. *Speakable and Unspeakable in Quantum Mechanics*. Cambridge: Cambridge University Press.

Bub, J., and R. Clifton. 1996. "A Uniqueness Theorem for Interpretations of Quantum Mechanics." *Studies in the History and Philosophy of Modern Physics* 27: 181–219.

Clifton, R. 1996. "The Properties of Modal Interpretations of Quantum Mechanics." *British Journal for the Philosophy of Science* 47: 371–98.

Deutsch, D. 1985. "Quantum Theory as a Universal Physical Theory." *International Journal of Theoretical Physics* 24: 1–41.

Dieks, D. 1989. "Quantum Mechanics without the Projection Postulate and Its Realistic Interpretation." *Foundations of Physics* 19: 1397–1423.

Everett, H. 1957. "Relative State Formulation of Quantum Mechanics." *Reviews of Modern Physics* 29: 454–62.

Geroch, R. 1984. "The Everett Interpretation." *Noûs* 18: 617–33.

Healey, R. 1989. *The Philosophy of Quantum Mechanics*. Cambridge: Cambridge University Press.

Kochen, S. 1985. "A New Interpretation of Quantum Mechanics." In P. Lahti and P. Mittelstaedt, eds., *Symposium on the Foundations of Modern Physics*. Singapore: World Scientific, 151–69.

Lockwood, M. 1989. *Mind, Brain, and the Quantum*. Oxford: Oxford University Press.

Putnam, H. 1975. "The Logic of Quantum Mechanics." *Philosophical Papers vol. 1*. Cambridge: Cambridge University Press, 174–97.

Stein, H. 1984. "The Everett Interpretation of Quantum Mechanics: Many Worlds or None?" *Noûs* 18: 635–52.