Appendix B. Knowings.

I shall flesh out here the idea that Nature is built out of knowings, not matter.

A typical knowing of the kind that quantum theory is built upon is a knowing that the pointer on the measuring device appears to lie between the numbers 6 and 7 on the dial. This is the sort of fact that all (or at least most) of science is built upon. It is quite complex. The idea that the appearance pertains to a dial on something that acts as a measuring device has a tremendous amount of education and training built into it. Yet somehow this knowing has this background idea built into it: that idea is a part of the experience.

William James (1910) says about perceptions:

"Your acquaintance with reality grows literally by buds or drops of perception. Intellectually and upon reflection you can divide these into components, but as immediately given they come totally or not at all."

This fits perfectly with Copenhagen quantum theory, which takes these gestalts as the basic elements of the theory. In the von Neumann/ Wigner type ontology adopted here there is, in association with this knowing, a collapse of the state vector of the universe. It is specified by acting on this state with a projection operator that acts on the degrees of freedom associated with the brain of the perceiver, and that reduces the state of the body/brain of the observer, and consequently also the state of the whole universe, to the part of that state that is compatible with this knowing.

So a human knowing is a complex experiential type of event that, however, according to the theory, occurs in conjunction with a correspondingly complex "physical" event that reduces the state of the the brain/body of the person to whom the experience belongs to the part of that state that is compatible with the knowing. [I shall use the word "physical" to denote the aspect of nature that is represented in the Hilbert-space description used in quantum theory: this aspect is the quantum analog of the physical description of classical physics.]

That "person" is a system consisting of a sequence of knowings bound together by a set of tendencies that are specified by the state of the universe. This state is essentially a compendium of prior knowings. However, these knowings are not merely human knowings, but more general events of which human knowings are a special case.

According to the strict Copenhagen interpretation, quantum theory is to be regarded as merely a set of rules for making predictions about human knowledge on the basis of prior human knowledge: horses and pigs do not make theoretical calculations using these ideas about operators in Hilbert space, and their "knowings" are not included in "our knowldge.

But in a science-based *ontology* it would be unreasonable to posit that human knowledge plays a singular role: human knowings must be assumed to be particular examples of a general kind of "knowings" that would include "horse knowings" and "pig knowings". These could be degraded in many ways compared to human knowings, and perhaps richer in some other dimensions, but they should still be of the same general ontological type. And there should have been some sort of things of this general ontological kind even before the emergence of life. [In the section, "What is Consciousness", I have tried to provide an intuition about what a knowing associated with a nonbiological system might be like.]

Science is an ongoing endeavor that is expected to develop ever more adequate (for human needs) ideas about the nature of ourselves and of the world in which we find ourselves. Newton himself seemed to understand this, although some of his successors did not. But the present stage of theoretical physics makes it clear that we certainly do not now know all the answers to even the most basic questions: physics is still very much in a groping stage when it comes to the details of the basic underlying structure. So it would be folly, from a scientific perspective, to say that we must give specific answers now to all questions, in the way that classical physics once presumed to do.

This lack of certainty is highlighted by the fact that the Copenhagen school could claim to give practical rules that worked in the realm of human knowledge without paying any attention to the question of how nonhuman knowings entered into nature. And no evidence contrary to Copenhagen quantum theory has been established. This lack of data about nonhuman knowledge would make it presumptuous, in a science-based approach, to try to spell out at this time details of the nature of nonhuman knowings, beyond the reasonable presumption that animals with bodies structurally similar to the bodies of human beings ought, to the extent they also behave like human beings, to have similar experiences. But knowings cannot be assumed to be always exactly the kinds of experiences that we human beings have: they could be quite different.

The knowings that I mentioned at the outset were percepts: knowings that appear to be knowings about things lying outside the person's body. But, according to the von Neummann/ Wigner interpretation, each such knowing is actually connected directly to the state of the person's body/brain, after that event has occurred. This state of the body/brain will, in the case of percepts of the external world, normally be correlated to aspects of the state of the universe that are not part of the body/brain. But experienced feelings, such as the feelings of warmth, joy, depression, devotion, patriotism, mathematical understandings, etc. are not essentially different from percepts: all are experiences that are associated with collapse events that reduce the state of the body/brain to the part of it that is compatible with the experience.

I have spoken here of a body/brain, and its connection to an experience. But what is this body/brain? It seems to be something different from the knowing that it is connected to. What is the nature of this connection?

The body/brain is an aspect of the quantum mechanically described state of the universe. This Hilbert-space state (sometimes called density matrix) is expressed as a complex-valued function of two vectors, each of which is defined over a product of spaces. Each such space corresponds to a degree of freedom of the universe. Any system is characterized by a certain set of degrees of freedom, and the state of that system is defined by taking the trace of the state of the universe over all other degrees of freedom, thereby eliminating from this state any explicit reference to those other degrees of freedom.

In this way the (quantum) state of each system is separately definable, and dependent only on its own degrees of freedom, even though the system itself is basically only an aspect of the whole universe. Each part (i.e., system) is separately definable, yet basically ontologically inseparable from the whole: that is the inescapable basic message of quantum theory. Each system has a state that depends only on its own degrees of freedom, and this system, as specified by its state, is causally pertinent, because each knowing is associated with some system, and the probabilities for its alternative possible knowings are specified by its own state, in spite of the fact that the system itself is fundamentally an inseparable part of the entire universe. It is the properties of the trace operation that reconciles these disparate requirements

The (quantum) state of the universe specifies only the probabilities for knowings to occur, and it generally undergoes an instantaneous global instantaneous jump when a new knowing occurs. But this probability, by virtue of the way it jumps when a new knowing occurs, and suddenly changes in regions far away from the system associated with the new knowing, and that it is formulated in terms of infinite sets of pssibilities that may never occur, is more like an idea or a thought than a material reality. Indeed, these properties of the state are exactly why the founders of quantum theory were led to the conclusion that the mathematical formalism that they created was about knowledge.

The state of the universe is the preserved compendium of all knowings. More precisely, it is an aspect of that compendium that expresses certain statistical properties pertaining to the next knowing. There is presumably some deeper structure, not captured by the properties expressed in the Hilbert-space mathematical structure, that fixes what actually happens.

The knowings that constitute our experiences are the comings into being of bits of knowledge, which join to form the knowledge that is represented by the (quantum) state of the universe. This gives an ontology based on knowings, with nothing resembling matter present. But the statistical causal structure of the sequence knowings is expressed in terms of equations that are analogs of the mathematical laws that governed the matter postulated to exist by the principles of classical mechanics. This connection to classical mechanics is enough to ensure a close similarity between the predictions of classical mechanics and those of quantum mechanics in many cases of interest, even though the two theories are based on very different mathematical structures.

———————————————————— I sent the above draft to some scientists and philosophers, and received a reply from the cognitive scientist Pat Hayes. In order to clarify, and enlarge upon, what has been said above, I conclude by giving some of Pat's comments and my replies to them.

[Italics indicates Pat's comment; Indented italics indicates his quote of a passage of mine.]

*My main problem is that I don't know what a 'knowing' is supposed to be. While it might not be reasonable to ask for a *definition* at this stage, nevertheless one needs some kind of guide for one's intuition in order to even follow what an author means when he uses a coined word or a technical usage like this. And you start with what is apparently a proto-typical example of a 'knowing':*

> *A typical knowing of the kind that quantum theory is built upon is a knowing that the pointer on the measuring device appears to lie between the numbers 6 and 7 on the dial.*

OK, I can grok that; so a 'knowing' is something like an event of perception, or a coming to believe something, or a realization that something is true, or a having of a feeling about something, etc.. But all of these kinds of examples clearly require there to be a knower who is experiencing these knowings. If these are the examples I am to use to understand the concept, then the idea of a pure knowing, bodiless and impersonal, seems simply incoherent.

I'm prepared to accept that pigs and dogs, and even maybe - for the sake of this argument - paramecia, might be said to have such knowings. But in all these cases there is a knower somewhere, which implies that until there was at least life of some kind, 'knowings' like this simply couldn't have existed.

Your own account seems to imply this also, even when it suddenly gets technical:

> *... In the von Neumann/ Wigner type ontology adopted here there is, in association with this knowing, a collapse of the state vector of the universe. It is specified by acting on this state with a projection operator that acts on the degrees of freedom associated with the brain of the perceiver, and that reduces the state of the body/brain of the observer, and consequently also the state of the whole universe, to the part of that state that is compatible with this knowing.*

Notice that the knowing is located in an observer; and in fact seems to be identified with a change of that observer's physical state. So (to labor the *point) if there were no observer, there would be no state to change, and hence no knowing.*

...

> *That "person" is a system consisting of a sequence of knowings bound together by a set of tendencies that are specified by the state of the universe. This state is essentially a compendium of prior knowings. However, these knowings are not merely human knowings, but more general events of which human knowings are a special case.*

But now you shuffled the cards. All the examples we have been given refer to human or human-like experience embodied in physical bodies. How do we generalise this intuition to these 'more general' events? What category is this new, more general, concept, of which we have no examples other than the proper subset with which we are all familiar? Can you give us just a sketch of what a "knowing" might have been at the time when, say, the universe's temperature was so high that atoms had not yet formed? Or inside a main-sequence star? Or in deep space where the 'boiling vacuum' is busily creating high-mass particle-antiparticle pairs and then reannihilating them so quickly that the violation of mass conservation doesn't escape the Heisenberg limit?

The theory of the early universe is not in very good shape: I think all experts agree that a lot of tinkering is still needed. The question of what role the quantum collapses played then has not really been put in. It is possible that the universe evolved for billions of years with no collapses until the first primitive life form occurred, and then there was a stupendous collapse. Perhaps the "anthropic principle" could be explained by some such scenario. But I think we should allow, at this point in the development of an adequate understanding of the early universe, also the possibility that there were collapses at earier times, prior to the emergence of life.

Perhaps there would be no "experiential-type" event associated with such a collapse. But there is the problem that there exists today no understanding of what it is that chooses the particular collapse event that actually occurs, insofar as one stays within the confines of the "physical" aspect of nature specified by the (Hilbert-space) state vector: that (Hilbert-space) aspect seems naturally suited to give only statistical aspects. Some other aspect of nature seems to be needed to fix what actually happens: i.e., to fix which one of the various possible collapses actually occurs. So it seems that a purely physical description of nature (i.e., Hilbert-space description) is not complete, ontologically: some other aspect seems to be needed.

This problem of the early universe is at least as hard as the mind-matter problem because of the sparsity of data. We have human brains in abundance, close at hand, and can do all sorts of experiments on them. So I do not think it makes much sense to say that we have to produce a detailed theory of the early universe as a part of a proposed approach to the mind-matter problem at the human level: things could go in many different ways in the extrapolation to the early universe.

But in answer to your challenge to what a "knowing" during the early stages of the universe "might have been" I refer you to section 14, "What is Consciousness?". The purpose of that section was to provide an answer to your question. If the collapses are basically to states defined by vibrations in certain degrees of freedom of the universe, then I can at least dimly imagine that the raw basic experience could be some basic prototype "sound" of those vibrations.

Of course, the "sounds" that we experience are surely tremendously enriched by the input elicited from our memories, and it is perhaps impossible was us to have a clear idea of the raw pure experience of the sound of a vibration outside of our human situation.

Still, if one wants to hold onto the idea that there are collapses outside animals, and indeed in the infant universe, and that *some* sort of generalization of human experience should be associated with such collapses, then I think that this idea that 'primordial experience is like the experience of the sounds of the vibrations that are being actualized' is an intuitively accessible notion.

It is a conceivable-to-me kind of primitive proto-experiential forerunner of our experiences, which could conceivably evolve into the kind of experience we know.

It would have theoretical ramifications, because my basic premise is the the only possible collapses are those that express the restriction imposed by a "knowing", *broadly defined* so as to include proto-experiences. Human experiences should be just a special complex kind of element in some class of elements that ought to have also simple primitive elements.

Since the experience of a "chord" is more complex than just the sum of the experiences of the individual notes there seems to be the possibility that all experiences could be basically like the experiencing of the sounds of those chords of vibrations that are being actualized.

These vibrations will be vibrations in certain degrees of freedom of the state of the universe, and hence in some physical system: I do not envisage any experience that is not tied to certain physical degrees of freedom, and to an actualization of some fact that is completely characterized by a collapse of the state vector of the universe: the collapse involves placing some restriction on those degrees of freedom.

It is undoubtedly true that "knowings" could be DEFINED so that they are emergent: so that that they just fade into nothingness as the system associated with the collapse loses various features that are present in normal human brains. One could define "knowings" narrowly in this way, and that might indeed be less abusive of language.

But IF human thoughts are somehow connected with the choice or selection of *which* (from among the various possible) collapses are made in human brains—and it is this connection what would give our thoughts some nonepiphenomenal role of the physical world—then it would seem that some *analog* of our thoughts would be needed for the other collapse. That is, if our thoughts are not epiphenomenal, but do play an important role in the determination of which state will be actualized, then it would seem that, for *collapse, some* analog of a human thought would be needed to play this causal role. That is my assumption, and is the basis of my broad way of defining "knowing".

The question, then, is how do the human knowings that we know become generalized when we try to expand the pragmatic/epistemological strict Copenhagen approach to an ontological one. The problems are indeed enormous, and that is why the Copenhagen people declined to get involved at any detailed level with the ontological issues. We have no data about the "knowings" of non-humans, and no data about collapses that are not associated with accretions in human knowledge. So within the bounds of contemporary science one is pretty much confined to the human mind-brain, plus the reasonable general principle that human mind-brain process should not be some unique process in nature: it *must* (I hold) be a special case of a more general process. So what I am saying here is: let us admit that human mind-brain process is just a special case of a general natural process, but nonetheless focus *for now* on the mind-brain human problem, where we have a kind of knowledge that we do not have in the general case. Science is to some extent opportunistic: it seizes upon opportunities as they open up; it focusses on issues where there is the possibility of relevant data.

> *... a probability is more like an idea or a thought than a material reality.*

That's a remarkable claim. I find it quite odd, in fact. After all, probabilities can often be measured objectively with arbitrary accuracy, given large enough populations. They seem like very objective, external, things; aspects of the universe rather than subjective, internal things like ideas or thoughts.

This is, of course, a very big issue for philosophers of science: Is probability to be thought of in the objective "propensity" sense (e.g., as recommended by Popper) or as a feature of what is known when someone knows less than everything that could be known. To make the idea of probability precise one usually gets involved with 'infinity', which may be more of an idea than an actual feature of nature, and with infinite imaginary ensembles. The large collection of future possible experiments to which the probabilistic predictions of quantum theory refer will, for the most part, never be actually performed.

*Maybe this is one explanation for our clash of intuitions. I see the probability waves in the Copenhagen interpretation as something real, external, objective. They are there even if nobody is measuring anything; they are what the world is in some sense really made of. Its a very odd picture compared the the classical one, but it doesnt require one to posit anything psychological or epistemic about them. They are *physical* probabilities (of a particle having a momentum or an energy, etc.), not epistemic probabilities.*

When Heisenberg describes "what really happens" (See chapter 3: The Vulgar Copenhagen Interpretation) he definitely uses this objective tendency (propensity) idea. But he does distinguish that from what the probability function of Copenhagen quantum theory represents, namely our knowledge, which suddenly jumps when our knowledge changes. Popper has written a strongly-worded denunciation of the Copenhagen interpretation (calling it 'The Great Quantum Muddle') in the book "Quantum Theory and Reality" (ed. Mario Bunge: Springer-Verlag). It is basically an attack on the use of the subjective interpretation by the Copenhagenites [a fact that he amply documents], and a polemic in favor of using, instead, the propensity interpretaion.

But my position on the nature of the quantum probability is intermediate. I do not adhere to the Copenhagen idea that the quantum formalism is MERELY a tool for scientists and engineers to make predictions about what their future experiences are likely to be. I ontologicalize, giving no special ontological role to human beings, and make the probability function objective: it is tied to an objective sort of knowledge to which all "observing systems" can contribute, be they human or not, or biological or not. The probability thus becomes objective, since the relevant knowledge is objective. Yet it is a representation of "knowledge", and it evolves as this knowledge grows, with a sudden jump associated with each new "observation", insofar as an observation means a new "knowing".

> *...And the fact that this probability suddenly changes when a new knowing occurs also makes the state more idea-like than matter-like.*

*But is that really true? Isn't it rather that the probability suddenly changes *for the observer*. But if someone else includes that observer into their external universe, then the first observer's observation makes no difference to his calculations.*

I agree that BESIDES the objective probabilities there can be subjective (Copenhagen-type) ones that reflect only the subset of knowings that a community of communicating human observers possesses and can use to make practical predictions about what their own future knowings are likely to be. This subjective probability associated with this communicating community of human being reflects only *their* knowledge, and changes only when *their* knowledge changes, whereas the objective probability changes when any observing system has a knowing.

Thus my approach allows the epistemological Copenhagen quantum theory [which is (weakly) objective in the sense that the probabilities do not depend upon any particular member of the community of communicating observers] to be naturally imbedded in a (strongly) objective ontological structure of the same general form.

The thing that needs to be explained, then, is why we all feel so vividly that what we perceive doesn't consist of probabilities but instead seems more like the classical world which we all once thought it actually was. Why do we have this mass illusion, as it were?

You ask why do our experiences have the form they have.

That is just what my paper is all about.

I think we do agree that your experiences are causally associated more directly with your own body/brain, than with some aspect of nature such as the record player out there that is playing the Verdi opera that you are listening to and hearing, or the performers and musicians that produced the music in the first place. There is a causal chain going back to them, but the aspect of nature most directly associated with your experiences is your own body/brain. I do not think that, in talking to you, I need to produce an argument for that.

But why do you hear the music rather than experiencing the patterns of neural firings in your body/brain directly as patterns of neural firings

If one follows my notion that your experience is the "sound" of vibrations in your brain, then all that is needed is a particular causal mapping that maps the acoustical vibrations in your ear to some appropriate corresponding neuronal vibrations in your brain, and you would be able to hear the music. Of course, there can be a tremendous accompaniment of other vibration brought about by associations. So our experience can be correspondingly complex.

But this possible vibration-sound association is subsidiary to my main argument.

The key premise of my theory that the each human experience is associated with a collapse that actualizes a template for action, of the body/brain, and that the experience is the experience of the intention of that action, which can be either to intend or to attend, where the latter is intention to update the body-world-mind schema, which is the brains representation of the body-world-mind. The "mind" part is an extention of the body-world part, and it includes our more abstract ideas. The fixing of an intention is represented in the physical realm [i.e., in the Hilbert-space representation of certain statistical properties of nature] by the actualization of a projected [i.e., intended] state of the body-world-mind schema. Natural selection is supposed to have coordinated the experience (i.e., the knowing) associated with the actualization of a body-world-mind schema to the mechanical evolution that follows from the causal structure of the body-brain in such a way that the caual effect of the actualization of a body-world-schema that is connected to a certain experienced intended action will, under appropriately directed attention, generally be to produce a subsequent state of the body/brain that is likely to have a collapse that actualizes a body-world-mind schema that corresponds to an experiencing of the intended action.

It seems to me that only in this way will our experiences be able to play the executive role in the control of our actions and thoughts that they appears to play. Natural selection ought then to hone and mold the connection between the experience and its effect, which is the collapse that brings the wave function of the body/brain—and more specifically the body-world-mind schema, which is the brains representation of the body-world-mind—into this causal accord with the later experience that follows the earlier experience that produced this effect.

This naturalistically generated connection between experiences, collapses, and templates for action formulated in terms of the body-world-mind schema explains the connection you required, namely the fact that our experiences seem to be direct knowings of the (perceived or intended) state of our "body", as we feel it, or the state of the "world" as it is classically conceived, or of ones "mind", or of all of these things in some combination.

The collapse is to a state of the brain that actualizes the brains representation of the body-world-mind. But you may ask: how do we get from the actualization of THE BRAIN'S REPRESENTATION of a body-world-mind to an experiencing, not of this representation by neural patterns of activity, but of this other presentation in terms of sounds, colors, perceptions of geometric forms, feelings of warmth and coldness, etc. Even given the presumption that the two representation are isomorphic, why do we experience the experiential qualities rather than their "physical representations in the brain"?

That is the question I have just answered.

But this answer hinges on understanding what is real and what is not.

In classical physics one imagines that the physical world is real, in and of itself, with no help or assistance from anyone's thoughts about it.

But in quantum theory it is the knowings that are the real things, and the knowings associated with human body/brains have the qualities that we human beings experience. These knowing are connected to each other in a way that can be codified, at least in some good approximation, and at least with regard to certain statistical regularities, in terms of a mathematical formalism that involves such notions as spacetime and Hilbert-space vectors, and that allows connections between our knowings to be represented mathematically, at least in some useful approximation, in terms of changes in a certain Hilbert-space vector that can be imagined to represent an objective state of knowledge that each knowing, human or otherwise, contributes to.

But just as the mathematical representation of the physical world provided by classical physics turned out to provide merely a partial characterization of certain properties of nature herself, so also should the Hilbert-space representation be understood as basically not the real thing itself, but rather a partial characterization that captures certain statistical PROPERTIES of the real thing. On the other hand, our human knowings are PARTS of the real thing.

Seems to be that I AM a body which has some experiences, not a experiential-type thing that uses a body as a 'host'.

What to you "seems to be" could be wrong and misleading: you cannot assume that what "seems to be" to you is necessarily what actually is.

A more secure assumption would be that the "seemings" occur; and that the content of these "seemings" includes, in its fringe, the impression that there is a bodily person, namely "Pat Hayes", that "has" these experiences. These experiential "seemings" are the "knowings" that I am talking about. And we agree that they are associated in some way with some appearance that goes by the name "Pat Hayes". The question is: What is the nature of these two appearances, the "seemings"and their apparent bodily homes, and how are these "seemings" related to "what seems to be".

> *In the light of the failure, at the basic level, of classical ideas about the nature of the physical world one cannot just take it for granted*

that the bodily aspect are basically in accord with ideas of classical mechanics,

This depends on scale. QM predicts that when observed at a 'human' scale, the universe will behave classically (to a reasonable approximation.)

Scale is not enough: there are meter-sized (and even kilometer-sized) systems that HUGELY fail to behave in accordance with classical ideas, but whose behaviour is described very accurately by QT. There is no evidence that functioning brains behave in accord with the classical concepts, as contrasted to quantum concepts: physicists have found many reasons, based on the more generally applicable principles of QT, why the classical principle could or should fail to work in living functioning brains.

Nobody feels that, say, oncology is fundamentally flawed because it uses classical physical models. If classical physics is good enough for medicine, why might it not be good enough for psychology?

Cells, like other physical systems, are certainly much better understood in ordinary physical terms than our thoughts are. We can understand the cell in terms of the idea that it IS a big collection of tiny particle that push and pull each other about: but we do not understand how, e.g., an experience of greeness IS a collection of such particles, or even is a happening composed of the activities of such classically conceived particles.

There is no evidence whatever for any kind of 'experience' existing without some complex physical substrate to be the experiencer. There is no evidence whatever, and plenty of evidence against, disembodied experience.

There is no real evidence even for dog-knowings. We do not have a secure enough theory of human knowing to extrapolate even to dogs: we are just going on what seems reasonable to scientifically oriented thinkers.

All the experiences that I am talking about are associated with physical systems, though not necessarily solid bodies. I do not know of any evidence against the possibility of the existence in association with liquid our gaseous physical systems of something that lies in some broad ontological class the contains also human experiences.

We shall have to wait and see whether collapses of the wave function are needed to account for the development of the early universe. But I have said that I doubt that there is a sharp boundary between life and nonlife. So one should expect collapses even without life.

But IF there are collapses in the infant universe, and IF human and animal knowings play an essential role in determining which collapses can occur in connection with human and animal bodies, then there ought to be some analog of human and animal knowings also for the early collapses: there ought to be something that plays a similar *causal role* in these early universe cases. How much experiential similarity they have to human knowings is very conjectural: I have stressed that the character of our experiences must be expected to be highly dependent on the structure of our brains: different kinds of brains should have different kinds of experience, and I have stressed that systems with no brains should therefore have nothing that is experientially very similar to a human experience.

It is by virtue of the *causal role* that they play that these events are analogous to human knowings. Yet if these events are causally similar it is reasonable to ask whether there is any *imaginable* way in which the primitive events "might conceivably" have some rudimentary or elemental proto-experiential aspect that could eventually develop, in conjunction with the development of the physical systems with which they are associated, into human-like experiences. I think the analogy to vibrations and music may suffice here.

I do not claim any "freedom" for consciousness: the choice or selection is not determined by the Hilbert-space description, because that is a representation merely of some statistical properties of nature. The failure of the actual choice to be described by a representation that describes only statistical properties does not entail any fundamental "freedom". Rather it opens the door to a deeper nonlocal causal structure based on the real things in nature, which include our knowings. [See my Appendix D]

I am careful to maintain fully the validity of the statistical predictions of quantum theory, with no biasing or other deviations in connection with consciousness. [See my section 11: Origin of the statistical rules]

I don't think that all of David Chalmer's arguments for the 'hard problem' being all that hard are convincing, but he does have a point: one does need to be able to at least find a connection principle remotely plausible. I can't imagine *what principles could connect experience to a "vibration in a physical degree of freedom". Just as a start, for example, what kind of psychological predictions would that make? What properties of awareness *in brains* would follow from this hypothesis?*

What is suggested is that every human experience is the" sound" of a complex chord of vibrations in certain degrees of freedom associated with a human brain. This would have empirical ramifications in neuro-psych. In mapping out the mind-brain connection, the relevant brain properties would be chords of vibrations of degrees of freedom of the brain. This is a very special sort of NCC (Neural Correlate of Consciousness).

To bring this idea into cognitive science, I believe one could design computer models in which certain chords of vibrations in certain elements are the central units of control in the processing. To make the system analogous to a quantum system there would be massive parallel processing that would generate alternative possible chords. When a chord reaches a "steady state" there would be a random choice, with weights controlled by prior performance success, that would either kill it off or permit it to activate a corresponding templated activity in the nonparallel processing part of the processor, and send a message to all parallel processing units that the new boundary condition is in place: the nonparallel processor is already committed to the execution of a certain template for action.

In the brain the quantum collapse automatically resets the boundary conditions globally and instantaneously without any need to send signals. The elements entering into the chord could be located in different parts of the brain without prior coordination of delicate timings: the coordination in timings could be fixed by the collapse, with quantum tunnelings and quantum diffusion serving to feed states that do not have the correct timings—and which, therefore, in a classical model would hang indefinitely in limbo—into nearby states that do have the correct timings. This feature of a quantum brain could give it a huge performance superiority over a classical brain. Hence I would expect natural selection to exploit it.

A key point in the quantum processing that I am referring to is that there is no need to maintain long-range coherence. Such coherence would seem to be very hard to do in a warm, wet brain.

To understand this last point in more detail, suppose the allowed experience involves a specific behaviour in each of ten elements scattered over the brain. and that each of the corresponding ten regions in the space of appropriate variables is bordered by nine other region that do not fit onto the allowed template, but which are equally likely to get populated, due to the effects of various (e.g., thermal) uncertainties. Classically there is only one chance in 10^{10} that the system will get into the good state that has the proper timings, and hence produce an experience and its associated collapse, rather than getting hung up in some non-productive state.

But a rapid quantum tunneling could quickly fill up the 'hole' (corresponding to the missing "good state") that would be present in the brain state that is left over after the good state is actualized. [The actualization of the the good state leaves a remainder, which is the original state minus the good state that was actualized, and hence was removed from this remainder.] Then that newly filled (by the quantum tunneling) good state could be actualized, etc. Thus all of the 10^{10} regions, only one of which is a good state (i.e., one that can be actualized by an experience) could be rapidly converted to actualized good states, whereas in the corresponding classical model only one of these 10^{10} states would lead to an experience, and the associated integrated action.

Quantum diffusion can have a similar effect. [See section 12: The Efficacy of Knowings]

> *Of course, the "sounds" that we experience are surely tremendously enriched by the input elicited from our memories, and it is perhaps impossible for us to have a clear idea of the raw pure experience of the sound of a vibration outside of our human situation.*

My problem is that I don't even accept that his idea makes sense. You just haven't given us enough to even begin to make judgements about what you are saying. To discuss difficulties of observation is premature.

OK. If the suggestion doesn't grab you, ignore it: it is not important that we have any intuitive idea of what the primordial proto-experiences are like. By the time the experiencing system has evolved to something as complex as a human being there is no compelling reason why the primordial experiential elements should still be accessible to our thoughts.

But I, in any case, can fathom the idea of elemental proto-experiences that are simpler than human experiences, yet that play a similar *causal role* in the determination of what collapses can occur.

> *But this possible vibration-sound association is subsidiary to my main argument.*
>
> *The key premise of my theory that the each human experience is associated with a collapse that actualizes a template for action, of the body/brain, and that the experience is the experience of the intention of that action, which can be either to intend or to attend, where the later is intention to update of the body-world-mind schema, which is the brains representation of the body-world-mind.*

This kind of vague template/schema-talk belongs in the 1950s. Cognitive science now has LIBRARIES of languages for 'schemas', analyses of their syntactic and semantic and computational properties, etc.. We know which ones admit sublinear search times on neural-net architectures using phase encoding, which are NP-complete (and therefore in general computationally intractable), how they are related to various kinds of model theory, etc. etc. etc.. One thing we know for sure is that most simple ideas of 'action schema' just aren't adequate to support human linguistic and perceptual performance. Now, what exactly do you mean by "schema" ? How are your schema related to Hilbert space?

See below.

> *The fixing of an intention is represented in the physical realm [i.e., in the Hilbert-space representation of certain statistical properties of nature] by the actualization of a projected [intended] state of the body-world-mind schema. Natural selection is supposed to have coordinated the experience (i.e., the knowing) associated with the actualization of a body-world-mind schema to the mechanical evolution that follows from the causal structure of the body-brain in such a way that the causal effect of the actualization of a body-world-schema that*

is connected to a certain experienced intended action will, under appropriately directed attention, generally be to produce a subsequent state of the body/brain that is likely to have a collapse that actualizes a body-world-mind schema that corresponds to an experiencing of the intended action.

That sounds like a cute idea, but how does this nonphysical natural selection work, exactly? Is it Darwinian? If so, how does a creature pass on its nonphysical traits to its offspring?

Yes, it is Darwinian. Suppose, just by way of example, that we accept the notion that an important real aspect of the universe is the existence of "sounds". These include not only the complex experiences that we human beings call sounds, but also much more primitive proto-experiences. Nature is composed of these "sounds" and a "causal" mathematical process that brings them into being in some sequence. This process has statistical regularities that are represented by a Hilbert-space state of the universe that changes each time a new chord occurs. So the experiential sounds are real happenings, and each such sound is associated with a set of degrees of freedom of this state, and in particular with the actualization of some brief "tone" (i.e., vibration) in that degree of freedom.

Just as a system of organ pipes (or a telephone tone-dialing system) can respond only to certain frequencies, so each physical system can respond only to certain sounds. Now the occurrence of an appropriate chord of sounds associated with a system activates a characteristic response, which might include producing in that system conditions that make it likely, according to the quantum statistical rules, that a similar chord will occur on that "instrument" slightly later. This cyclic behaviour would give the system a certain tendency to endure.

A system that cycles back to itself physically tends to endure, in the sense that its physical (i.e., Hilbert-space) form regenerates itself.

The effect of the actualization associated with the occurrence of a chord will have effects on the system. Since the information put into the system by the actualizatization has a certain mathematical structure it is possible that some of the external configurations of that system produced by the actualization of the chord can exhibit a certain isomorphism to the structure of the chord.

In a biological system both the structure of the allowed chords and the connection between the chords and the resulting bodily configuration will be determined by the physical structure of the system, which is controlled in part by the genes.

Suppose the living system evolves in such a way that the actualization of certain chords activate motor systems in a way such that 1) motions of the system are generated that 2) activate sensing systems that 3) create the likelihood that a chord similar to the original chord will occur slightly later. It is a key part of my theory (following James) that each experience, though experienced as a whole, has a sequence of temporal slices corresponding to parts that enter the sequence of experiences as temporally displaced components. [cf. section 8. Ontology: Quantum Stuff.] This allows aspects of the experience to be compared, and progressing activities to be monitored, within the slowly changing fringe of the experience.

In animal-type systems the progression of events is not rigidly fixed. The classical and quantum fluctuations allow branchings of the temporal dynamical development to occur: the animal can do either this or that.

In our own case, the sequence of conscious events that control/curtail the proliferation of quantum possibilities is such that the sequence of events can monitor and control the activities of the body/brain system. That is, there is this feed-back that I have been describing where the activity generated by the experiential event is such as to generate (under proper conditions of attention) a similar follow-up experiential event. In this situation the sequence of events can function to keep the quantum generation of successive possibilities on a coherent track: selections between coherent courses of action will be made, but there will not be pandemonium. The strong tendency will be for successive quantum choices to be between actions that are generally in line with what has been chosen just before.

For systems strongly controlled by an associated sequence of experiential events there must be such a recurrence type of structure at the experiential level in order for directed behaviour to be possible. If the animal system were not organized with this sort of recurrence, then the successive selections needed to choose between the quantum possibilities that are being continually generated would produce chaotic activity, rather than activity that has inertia: i.e., a sustained direction. The possibility of sustained directed activity is, of course, essential to survival.

So Darwian evolution should produce animals having this crucial connection between the experiences (which are the cause of the collapses of the wave function that are trimming away the quantum proliferation of alternative possible branches of temporal development) and the effects of these collapses both on body/brain activity and on the consequent follow-up experiences.

Because each experience has within itself this sequence of temporal slices, it becomes possible for the experience to include a sense of direction, and intention. That is, the experience can include not only the temporal slices (all existing as components of the one present thought, as specified by James) but also a slowly changing fringe part that surveys the structure of the passing parade of slices. This complex structure could include also a "projected" image of an expected slice.

> *It seems to me that only in this way will our experiences be able to play the executive role in the control of our actions and thoughts that they appears to play.*

Seems to me to much more likely that this 'appear' is simply an illusion, and that our conscious experiences don't, as a matter of fact, play such a simple 'executive' role, but are more like a story we tell ourselves as we go on to make sense of our own actions.

It may be significant that if one adheres to what contemporary physics tells us about our brains, rather than starting from a fundamentally incorrect conception of the nature of the brain, then we do not have to invent any tales of illusions and deception: we can preserve our deepest intuitions about our own nature by simply following through on what science has revealed about ourselves. Tales of delusion and self deception might be acceptable if they allowed us to bringing our ideas INTO accord with science, but not for trying make sense of a conception about our natures that science has found to be profoundly wrong.

> *Natural selection ought then to hone and mold the connection between the experience and its effect, which is the collapse that brings*

the wave function of the body/brain—and more specifically the body-world-mind schema, which is the brains representation of the body-world-mind—into this causal accord with the subsequently occurring experience.

This naturalistically generated connection between experiences, collapses, and templates for action formulated in terms of the body-world-mind schema explains the connection you required, namely the fact that our experiences seem to be direct knowings of the (perceived or intended) state of our "body", as we feel it, or the state of the "world" as it is classically conceived, or of one's "mind", or of these things in some combination.

Without a LOT more detail, this account doesnt explain anything at all.

This account is indeed only a general overview.

OK, let me challenge you directly. Once, I watched my mother break an egg against the side of a bowl, open the shell carefully and drop the contents into the bowl, all with one hand. I was astonished and asked her to show me how to do it, which she did, and I've been able to do it ever since. Now, this seems to me about as clear a case as one can come up with of a 'knowing'; I came to know something (both intellectually and somatically) I had not previously known. Right now I am involved in a group effort to formalise at least part of this knowledge in a way that a robot might find useful; it turns out to be surprisingly complicated. How can this 'knowing' be represented mathematically in terms of changes in a Hilbert-space vector?

That is a fair question.

Of course, it is supposed to be formulated in the context of some appropriate program that is simulating an experiencing quantum robot. The experiencings/knowings would be represented, let us say, by chords of sounds. If we allow, say N "tones" and M intensities for each tone we a space of possible experiences/knowings that has M^N distinct elements. The body/brain has set of D degrees of freedom, with D>>N: the N tones are associated with some influential subset of N of the D degrees of freedom of the body/brain of the experiencing/knowing robot. The robot has a non-parallel processing part that provides the sensing channels and the motor channels, and that corresponds to the firing-rate-controlled processes in animals, and a parallel processing part that corresponds to the essential quantum processes that are associated with knowings. This latter part operates within the boundary conditions specified by the non-parallel processing part. It generates the optional courses of possible actions: these options are trimmed back when an experience/knowing occurs. The various processes are governed by deterministic equations of motion, but each knowing is specified also by a tie into a psuedo-random number generator, so that the system simulates a quantum system.

We are to suppose that before the egg-breaking learning occurs, the system already "knows" about the appearances and handling of eggs and bowls. This means that cycles have been establish (by prior learning) that connect various "experiences" pertaining to the simultaneous seeing, and feeling, and handling of eggs into feed-back loops of the kind I have described, wherein an experience (represented in this simulation or mock-up by a fixing of certain of the N degrees of freedom of the body/brain) intiates a motor sequence that intitiates a sensing sequence that produces a high likelihood of producing in a later temporal slice in the knowing a likeness of the aspect of the experience that initiated the experience.

If your learning involved hearing your Mother's explanation of what she was doing, them we must expand the model to include language skills, and motor-visual-tactile interpretations of heard instructions. I believe this same sort of model can cover language skills, provided the space of possible knowings is large enough, and the appropriate linkages to motor and sensor processes exist: I emphasize that each 'experience' contains a whole sequence of temporal slices, and that actions controlled by such whole 'experiences' have been honed by learning so that the actions generate reactions on the evolving sequence of slices such that the generated incoming slices are similar to, and temporal extensions of, the sequence that generated them. I look upon learning as the generation of successful repertoires of experiences, each containing a sequence of temporal slices, that generate smooth prolongations of these sequences.

Sequences that produce smooth prolongations of themselves will recur and thus reinforce themselves by neural mechanisms in animal brains, or by some corresponding designed feature of the robot: such sequences will be "facilitated", and will tend to arise more easily.

Of course, chords have subchords, so there are vast possibilities of combinations of already established subchords. There must be, honed into brain structure by evolution and training, or by designed features in the robot, variable rules of composition that tend to extend recognition to combinations of subchords that often arrive together, or in some particular sequence.

I do not see the question of exactly how the designed facilitations work as particularly challenging.

Your learning how to crack an egg on a bowl with one hand is an example of fitting together a number of already learned repertoires, about handling such things as eggs, by running through the sequence a few times either with actual motions generated by combining already known motions in a new way, or perhaps by running through the sequence only in your imagination, which would involve brain action, but no appreciable bodily motion.

This answer to your challenge may be far less detailed than you would like, but I believe that it is enough of a blueprint to allow me to write a C++ program that would function satifactorily: combinations of already established objects, representing experiences and their actions and reactions back onto the class of experiences, would be the building blocks, with each experience containing a sequence of temporal slices, and causing a physical action. These 'experiences' could act nonlocally on many degrees of freedom simultaneously. These potentially very complex controlling elements, each with sequences of temporal slices, and an overarching superstructure that gives structure to the whole sequence, and with a tendency to regenerate themselves, would be the computer analogs of human knowings.