

Quantum mechanics as quantum information, mostly

Christopher A. Fuchs

To cite this article: Christopher A. Fuchs (2003) Quantum mechanics as quantum information, mostly, Journal of Modern Optics, 50:6-7, 987-1023

To link to this article: <http://dx.doi.org/10.1080/09500340308234548>



Published online: 03 Jul 2009.



Submit your article to this journal [↗](#)



Article views: 107



View related articles [↗](#)



Citing articles: 35 View citing articles [↗](#)

Quantum mechanics as quantum information, mostly†

CHRISTOPHER A. FUCHS

Bell Labs, Lucent Technologies, 600–700 Mountain Avenue,
Room 1D-236, Murray Hill, NJ 07974, USA

(Received 9 September 2002; revision received 17 December 2002)

Abstract. In this paper, I try to cause some good-natured trouble. The issue is, when will we ever stop burdening the taxpayer with conferences devoted to the quantum foundations? The suspicion is expressed that no end will be in sight until a means is found to reduce quantum theory to two or three statements of crisp physical (rather than abstract, axiomatic) significance. In this regard, no tool appears better calibrated for a direct assault than quantum information theory. Far from a strained application of the latest fad to a time-honoured problem, this method holds promise precisely because a large part—but not all—of the structure of quantum theory has always concerned information. It is just that the physics community needs reminding.

1. Introduction

Quantum theory as a weather-sturdy structure has been with us for 75 years. Yet, there is a sense in which the struggle for its construction remains. I say this because not a year has gone by in the last 30 when there has not been a conference devoted to some aspect of quantum foundations‡.

How did this come about? What is the cause of this year-after-year sacrifice to the ‘great mystery’? Whatever it is, it cannot be for want of a self-ordained solution: go to any meeting, and it is like being in a holy city in great tumult. You will find all the religions with all their priests pitted in holy war—the Bohmians [2], the Consistent Historians [3], the Transactionalists [4], the Spontaneous Collapseans [5], the Einselectionists [6], the Contextual Objectivists [7], the outright Everettics [8], and many more beyond that. They all declare to see the light, the ultimate light. Each tells us that if we will accept their solution as our savior, then we too will see the light.

But there has to be something wrong with this! If any of these priests had truly shown the light, there simply would not be the year-after-year conference. The verdict seems clear enough: if we—i.e. the set of people who might be reading this paper—really care about quantum foundations, then it behooves us to ask why these meetings are happening and find a way to put them to a stop.

My view of the problem is this. Despite the accusations of incompleteness, non-sensicality, irrelevance, and surreality one often sees one religion making

† This paper, though containing some new material in sections 3 and 5, is predominantly a précis of [1].

‡ See [1] for a table of 30 explicit examples.

against the other, I see little to no difference in *any* of their canons. They all look equally detached from the world of quantum practice to me. For, though each seems to want a firm reality within the theory—i.e. a single God they can point to and declare, ‘There, that term is what is real in the universe even when there are no physicists about’—none have worked very hard to get out of the Platonic realm of pure mathematics to find it.

What I mean by this deliberately provocative statement is that in spite of the differences in what the churches label to be ‘real’ in quantum theory, they nonetheless all proceed from the same *abstract* starting point—the standard textbook accounts of the *axioms* of quantum theory.

‘But what nonsense is this’, you must be asking. ‘Where else could they start?’ The main issue is this, and no one has said it more clearly than Rovelli [9]. Where present-day quantum-foundation studies have stagnated in the stream of history is not so unlike where the physics of length contraction and time dilation stood before Einstein’s 1905 paper on special relativity.

The Canon for Most of the Quantum Churches: The Axioms (plain and simple)

1. For every system, there is a complex Hilbert space \mathcal{H} .
2. States of the system correspond to projection operators onto \mathcal{H} .
3. Those things that are observable *somehow* correspond to the eigenprojectors of Hermitian operators.
4. Isolated systems evolve according to the Schrödinger equation.

⋮

The Lorentz transformations have the name they do, rather than, say, the Einstein transformations, for a good reason: Lorentz had published some of them as early as 1895. Indeed one could say that most of the empirical predictions of special relativity were in place well before Einstein came onto the scene. But that was of little consolation to the pre-Einsteinian physics community striving so hard to make sense of electromagnetic phenomena and the luminiferous ether. Precisely because the *only* justification for the Lorentz transformations appeared to be their *empirical adequacy*, they remained a mystery to be conquered. More particularly, this was a mystery that heaping further *ad hoc* (mathematical) structure onto could not possibly solve.

What was being begged for in the years between 1895 and 1905 was an understanding of the *origin* of that abstract, mathematical structure—some simple, crisp *physical* statements with respect to which the necessity of the mathematics would be indisputable. Einstein supplied that and became one of the greatest physicists of all time. He reduced the mysterious structure of the Lorentz transformations to two simple statements expressible in common language: (1) the speed of light in empty space is independent of the speed of its source and (2) physics should appear the same in all inertial reference frames. The deep significance of this for the quantum problem should stand up and speak overpoweringly to anyone who admires these principles.

Einstein’s move effectively stopped all further debate on the origins of the Lorentz transformations. Outside of the time of the Nazi regime in Germany, I suspect there have been less than a handful of conferences devoted to

'interpreting' them. Most importantly, with the supreme simplicity of Einstein's principles, physics became ready for 'the next step'. Is it possible to imagine that any mind—even Einstein's—could have made the leap to general relativity directly from the original, abstract structure of the Lorentz transformations? A structure that was only empirically adequate? I would say no. Indeed, one can dream of the wonders we will find in pursuing the same strategy of simplification for the quantum foundations.

The task is not to make sense of the quantum axioms by heaping more structure, more definitions, more science-fiction imagery on top of them, but to throw them away wholesale and start afresh. We should be relentless in asking ourselves: from what deep *physical* principles might we *derive* this exquisite mathematical structure? Those principles should be crisp; they should be compelling. They should stir the soul. When I was in junior high school, I sat down with Martin Gardner's book *Relativity for the Million* and came away with an understanding of the subject that sustains me today: the concepts were strange, but they were clear enough that I could get a grasp on them knowing little more mathematics than simple arithmetic. One should expect no less for a proper foundation to quantum theory. Until we can explain quantum theory's *essence* to a junior high-school or high-school student and have them walk away with a deep, lasting memory, we will not have understood a thing about the quantum foundations.

Symbolically, where we are:	Where we need to be:
$x' = \frac{x-vt}{(1-v^2/c^2)^{1/2}}$	Speed of light is constant.
$t' = \frac{t-vx/c^2}{(1-v^2/c^2)^{1/2}}$	Physics is the same in all inertial frames.

So, throw the existing axioms of quantum mechanics away and start afresh! But how to proceed? I myself see no alternative but to contemplate deep and hard the tasks, techniques and implications of quantum information theory. The reason is simple and I think inescapable. Quantum mechanics has always been about information; it is just that the physics community has forgotten this.

This, I see as the line of attack we should pursue with relentless consistency: the quantum system represents something real and independent of us; the quantum state represents a collection of subjective degrees of belief about *something* to do with that system (even if only in connection with our experimental kicks to it). The structure called quantum mechanics is about the interplay of these two things—the subjective and the objective. The task before us is to separate the wheat from the chaff. If the quantum state represents subjective information, then how much of its mathematical support structure might be of that same character? Some of it, maybe most of it, but surely not all of it.

Our foremost task should be to go to each and every axiom of quantum theory and give it an information theoretic justification if we can. Only when we are finished picking off all the terms (or combinations of terms) that can be interpreted as subjective information will we be in a position to make real progress in quantum foundations. The raw distillate left behind—miniscule though it may be with

respect to the full-blown theory—will be our first glimpse of what quantum mechanics is trying to tell us about nature itself.

Let me try to give a better way to think about this by making use of Einstein again. What might have been his greatest achievement in building general relativity? I would say it was in his recognizing that the ‘gravitational field’ one feels in an accelerating elevator is a coordinate effect. That is, the ‘field’ in that case is something induced purely with respect to the description of an observer. From this view, the programme of trying to develop general relativity boiled down to recognizing all the things within gravitational and motional phenomena that should be viewed as consequences of our coordinate choices. It was in identifying all the things that are ‘numerically additional’ to the observer-free situation—i.e. those things that come about purely by bringing the observer (scientific agent, coordinate system, etc.) back into the picture.

Quantum Mechanics: *The Axioms and Our Imperative!*

States correspond to density operators ρ over a Hilbert space \mathcal{H} .

Give an information theoretic reason if possible!

Measurements correspond to positive operator-valued measures (POVMs) $\{E_d\}$ on \mathcal{H} .

Give an information theoretic reason if possible!

\mathcal{H} is a complex vector space, not a real vector space, not a quaternionic module.

Give an information theoretic reason if possible!

Systems combine according to the tensor product of their separate vector spaces, $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$.

Give an information theoretic reason if possible!

Between measurements, states evolve according to trace-preserving completely positive linear maps.

Give an information theoretic reason if possible!

By way of measurement, states evolve (up to normalization) via outcome-dependent completely positive linear maps.

Give an information theoretic reason if possible!

Probabilities for the outcomes of a measurement obey the Born rule for POVMs $\text{tr}(\rho E_d)$.

Give an information theoretic reason if possible!

The distillate that remains—the piece of quantum theory with no information theoretic significance—will be our first unadorned glimpse of ‘quantum reality’. Far from being the end of the journey, placing this conception of nature in open view will be the beginning of a new physics.

This was a breakthrough. For in weeding out all the things that can be interpreted as coordinate effects, the fruit left behind becomes clear to sight: it is the Riemannian manifold we call space–time—a mathematical object, the study of which one hopes will tell us something about nature itself, not merely about the observer in nature.

The dream I see for quantum mechanics is just this. Weed out all the terms that have to do with gambling commitments, information, knowledge and belief, and what is left behind will play the role of Einstein’s manifold. That is our goal. When we find it, it may be little more than a miniscule part of quantum *theory*.

But being a clear window into nature, we may start to see sights through it we could hardly have imagined before†.

2. Summary

This paper is about taking the Introduction's imperative seriously, although it contributes only a small amount to the labour it asks. Just as in the founding of quantum mechanics, this is not something that will spring forth from a single mind sheltered in a medieval college. It is a task for a community with diverse but productive points of view. The quantum information community is nothing if not that [11]. 'Philosophy is too important to be left to the philosophers', John Archibald Wheeler once said. Likewise, I am apt to say for the quantum foundations.

The structure of the remainder of the paper is as follows. In section 3, I reiterate the cleanest argument I know that the quantum state is solely an expression of subjective information—the information one has about a quantum system. It has no objective reality in and of itself‡. The argument is then refined by considering the phenomenon of quantum teleportation [14].

In section 4, entitled 'Information about what?', I tackle that very question head-on. The answer is 'the potential consequences of our experimental interventions into nature'. Once freed from the notion that quantum measurement ought to be about revealing traces of some pre-existing property (or beable [15]), one finds no particular reason to take the standard account of measurement (in terms of complete sets of orthogonal projection operators) as a basic notion.

† I should point out, however, that in contrast to the picture of general relativity, where reintroducing the coordinate system—i.e. reintroducing the observer—changes nothing about the manifold (it only tells us what kind of sensations the observer will pick up), I do not suspect the same for the quantum world. Here I suspect that reintroducing the observer will be more like introducing matter into pure space-time, rather than simply gridding it off with a coordinate system. 'Matter tells space-time how to curve (when matter is there), and space-time tells matter how to move (when matter is there)' [10]. Observers, scientific agents, a necessary part of reality? No. But do they tend to change things once they are on the scene? Yes. If quantum mechanics can tell us something deep about nature, I think it is this.

‡ Previously, I have not emphasized so much the 'radical' Bayesian idea that the probability one *ascribes* to a phenomenon amounts to *nothing* other than the gambling commitments one is willing to make on it. To the radical Bayesian, probabilities are subjective all the way to the bone. Here, I try to turn my earlier *de-emphasis* around. In particular, because of the objective overtones of the word 'knowledge'—i.e. that a particular piece of knowledge is either 'right' or 'wrong'—I try to steer clear of the term as much as possible. The conception lurking in the background of this paper is that there is simply no such thing as a 'right and true' quantum state. In all cases, a quantum state is specifically and only a mathematical symbol for capturing a set of beliefs or gambling commitments. Thus I variously call quantum states 'beliefs', 'states of belief', 'information' (though, by this I mean 'information' in a more subjective sense than is common in the quantum information community), 'judgements', 'opinions', and 'gambling commitments'. Believe me, I understand fully well the jaws that will drop from the adoption of this terminology. However, if the reader finds that this gives him a sense of butterflies in the stomach—or fears that I will become a solipsist [12] or a crystal-toting New Age practitioner of homeopathic medicine [13]—I hope he will keep in mind that this attempt to be absolutely frank about the subjectivity of *some* of the terms in quantum theory is part of a larger programme to delimit the terms that *can be* interpreted as objective in a fruitful way.

Indeed quantum information theory, with its emphasis on the utility of generalized measurements or positive operator-valued measures (POVMs) [16], suggests one should take those entities as the basic notion instead. The productivity of this point of view is demonstrated by the enticingly simple Gleason-like derivation of the quantum probability rule recently found by Busch [17] and, independently, by Renes and collaborators [18]. Contrary to Gleason's original theorem [19], this theorem works just as well for two-dimensional Hilbert spaces, and even for Hilbert spaces over the field of rational numbers. In section 4.1, I start the process of defining what it means—from the Bayesian point of view—to accept quantum mechanics as a theory. This leads to the notion of fixing a fiducial or standard quantum measurement for defining the very meaning of a quantum state.

In section 5, I ask whether entanglement is all it is touted to be as far as quantum foundations are concerned. That is, is entanglement really as Schrödinger said, 'the characteristic trait of quantum mechanics, the one that enforces its entire departure from classical lines of thought?' To combat this, I give a simple derivation of the tensor-product rule for combining Hilbert spaces of individual systems which takes the structure of *localized* quantum measurements as its starting point. In particular, the derivation makes use of Gleason-like considerations in the presence of classical communication. With the tensor-product structure established, the very notion of entanglement follows in step. This shows how entanglement, just like the standard probability rule, is secondary to the structure of quantum measurements. Moreover, 'locality' is built in at the outset; there is simply nothing mysterious and non-local about entanglement.

In section 6, I ask why one should expect the Bayes' rule for updating quantum state assignments upon the completion of a measurement to take the form it actually does. Along the way, I give a simple derivation that one's information always increases on average for *any* quantum mechanical measurement that does not itself discard information. (Despite the appearance otherwise, this is not a tautology!) Most importantly, the proof technique used for showing the theorem indicates an extremely strong analogy between quantum collapse and Bayes' rule in classical probability theory: up to an overall unitary 'readjustment' of one's final probabilistic beliefs—the readjustment takes into account one's initial state for the system as well as one's description of the measurement interaction—quantum collapse is *precisely* Bayesian conditionalization. In section 6.1, I complete the process started in section 4.1 and describe quantum measurement in Bayesian terms: an everyday measurement is any *I-know-not-what* that leads to an application of Bayes' rule.

In section 7, I argue that, to the extent that a quantum state is a subjective quantity, so must be the assignment of a state-change rule $\rho \rightarrow \rho_d$ for describing what happens to an initial quantum state upon the completion of a measurement—generally some POVM—whose outcome is d . In fact, the levels of subjectivity for the state and the state-change rule must be precisely the same for consistency's sake. To draw an analogy to Bayesian probability theory, the initial state ρ plays the role of an *a priori* probability distribution $P(h)$ for some hypothesis, the final state ρ_d plays the role of a posterior probability distribution $P(h|d)$, and the state-change rule $\rho \rightarrow \rho_d$ plays the role of the 'statistical model' $P(d|h)$ enacting the transition $P(h) \rightarrow P(h|d)$. To the extent that *all* Bayesian probabilities are subjective—even the probabilities $P(d|h)$ of a statistical model—so is the mapping $\rho \rightarrow \rho_d$. Specializing to the case that no information is gathered, one finds that the

trace-preserving completely positive maps that describe quantum time-evolution are themselves nothing more than subjective judgements.

In section 8, I review the parts of quantum mechanics argued to be subjective in character and reiterate that such an analysis cannot be the end of the journey.

Finally, in section 9, I flirt with the most tantalizing question of all: why the quantum? There is no answer here, but I do not discount that we are on the brink of finding one. In this regard no platform seems firmer for the leap than the very existence of quantum cryptography and quantum computing. The world is sensitive to our touch. It has a kind of ‘Zing!’ that makes it fly off in ways that were not imaginable classically. The whole structure of quantum mechanics—*it is speculated*—may be nothing more than the optimal method of reasoning and processing information in the light of such a fundamental (wonderful) sensitivity. As a concrete proposal for a potential mathematical expression of ‘Zing!’, I consider the integer parameter D traditionally ascribed to a quantum system by way of its Hilbert-space dimension.

3. Why information?

Einstein was the master of clear thought; I have expressed my opinion on this for both special and general relativity. But I can go further. I would say he possessed the same great penetrating power when it came to analysing the quantum too. For even there, he was immaculately clear and concise in his expression. In particular, he was the first person to say in unambiguous terms why the quantum state should be viewed as information (or, to say the same thing, as a representation of one’s beliefs and gambling commitments, credible or otherwise).

His argument was simply that a quantum-state assignment for a system can be forced to go one way or the other by interacting with a part of the world that should have no causal connection with the system of interest. The paradigm here is the one well known through the Einstein–Podolsky–Rosen paper [20], but simpler versions of it had a long pre-history with Einstein [21] alone.

The best was in essence this. Take two spatially separated systems A and B prepared in some entangled quantum state $|\psi^{AB}\rangle$. By measuring one or another of two observables on system A alone, one can *immediately* write down a new state for system B. Either the state will be drawn from one set of states $\{|\phi_i^B\rangle\}$ or another $\{|\eta_i^B\rangle\}$, depending upon which observable is measured. The key point is that it does not matter how distant the two systems are from each other, what sort of medium they might be immersed in, or any of the other fine details of the world. Einstein concluded that whatever these things called quantum states *be*, they cannot be ‘real states of affairs’ for system B alone. For, whatever the real, objective state of affairs at B is, it should not depend upon the measurements one makes on a causally unconnected system A.

Thus one must take it seriously that the new state (either a $|\phi_i^B\rangle$ or a $|\eta_i^B\rangle$) represents information about system B. In making a measurement on A, one learns something about B, but that is where the story ends. The state change cannot be construed to be something more physical than that. More particularly, the final quantum state for B cannot be viewed as more than a reflection of some tricky combination of one’s initial information and the knowledge gained through the measurement. Expressed in the language of Einstein, the quantum state cannot be a ‘complete’ description of the quantum system.

Here is the way Einstein put it to Michele Besso in a 1952 letter [22]:

What relation is there between the 'state' ('quantum state') described by a function ψ and a real deterministic situation (that we call the 'real state')? Does the quantum state characterize completely (1) or only incompletely (2) a real state?...

I reject (1) because it obliges us to admit that there is a rigid connection between parts of the system separated from each other in space in an arbitrary way (instantaneous action at a distance, which doesn't diminish when the distance increases). Here is the demonstration: [*The argument he uses is the same as the one reported above*].

If one considers the method of the present quantum theory as being in principle definitive, that amounts to renouncing a complete description of real states. One could justify this renunciation if one assumes that there is no law for real states—i.e. that their description would be useless. Otherwise said, that would mean: laws don't apply to things, but only to what observation teaches us about them. (The laws that relate to the temporal succession of this partial knowledge are however entirely deterministic.)

Now, I can't accept that. I think that the statistical character of the present theory is simply conditioned by the choice of an incomplete description.

There are two issues in this letter worth disentangling. (1) Rejecting the rigid connection of all nature—that is to say, admitting that the very notion of *separate systems* has any meaning at all—one is led to the conclusion that a quantum state cannot be a complete specification of a system. It must be information, at least in part. This point should be placed in contrast to the other well-known facet of Einstein's thought: namely, (2) an unwillingness to accept such an 'incompleteness' as a necessary trait of the physical world.

It is quite important to recognize that the first issue does not entail the second. Einstein had that firmly in mind, but he wanted more. His reason for going the further step was, I think, well justified *at the time* [23]:

There exists ... a simple psychological reason for the fact that this most nearly obvious interpretation is being shunned. For if the statistical quantum theory does not pretend to describe the individual system (and its development in time) completely, it appears unavoidable to look elsewhere for a complete description of the individual system; in doing so it would be clear from the very beginning that the elements of such a description are not contained within the conceptual scheme of the statistical quantum theory. With this one would admit that, in principle, this scheme could not serve as the basis of theoretical physics.

But the world has seen much in the meantime. The last 19 years have given confirmation after confirmation that the Bell inequality (and several variations of it) are indeed violated by the physical world. The Kochen–Specker no-go theorems have been meticulously clarified to the point where simple textbook pictures can be

drawn of them [24]. Incompleteness, it seems, is here to stay: the theory prescribes that no matter how much we know about a quantum system—even when we have *maximal* information about it—there will always be a statistical residue. There will always be questions that we can ask of a system for which we cannot predict the outcomes. In quantum theory, maximal information is simply not complete information [25]. But neither can it be completed. As Wolfgang Pauli once wrote to Markus Fierz [26], ‘The well-known ‘incompleteness’ of quantum mechanics (Einstein) is certainly an existent fact somehow-somewhere, but certainly cannot be removed by reverting to classical field physics.’ Nor, I would add, will the mystery of that ‘existent fact’ be removed by attempting to give the quantum state an ontological status.

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

Indeed I think this is one of the deepest questions we can ask and still hope to answer. But first things first. The more immediate question for anyone who has come this far—and one that deserves to be answered forthright—is what is this information symbolized by a $|\psi\rangle$ actually about? I have hinted that I would not dare say that it is about some kind of hidden variable (as the Bohmian might) or even about our place within the universal wavefunction (as the Everettic might).

Perhaps the best way to build up to an answer is to be true to the theme of this paper. Let us forage the phenomena of quantum information to see if we might first refine Einstein’s argument. One need look no further than to the phenomenon of quantum teleportation [14]. Not only can a quantum-state assignment for a system be forced to go one way or the other by interacting with another part of the world of no causal significance, but, for the cost of two bits, one can make that quantum state assignment anything one wants it to be.

Such an experiment starts out with Alice and Bob sharing a maximally entangled pair of qubits in the state $|\psi^{AB}\rangle = |0\rangle|0\rangle + |1\rangle|1\rangle$. Bob then goes to any place in the universe he wishes. Alice in her laboratory prepares another qubit with any state $|\psi\rangle$ which she ultimately wants to impart onto Bob’s system. She performs a Bell-basis measurement on the two qubits in her possession. In the same vein as Einstein’s thought experiment, Bob’s system immediately takes on the character of one of the states $|\psi\rangle$, $\sigma_x|\psi\rangle$, $\sigma_y|\psi\rangle$, or $\sigma_z|\psi\rangle$. But that is only insofar as Alice is concerned†. Since there is no (reasonable) causal connection between Alice and Bob, it must be that these states represent the possibilities for Alice’s updated beliefs about Bob’s system.

If now Alice broadcasts the result of her measurement to the world, Bob may complete the teleportation protocol by performing one of the four Pauli rotations (I , σ_x , σ_y , σ_z) on his system, conditioning it on the information he receives.

†As far as Bob is concerned, nothing whatsoever changes about the system in his possession: it started in the completely mixed state $\rho = \frac{1}{2}I$ and remains that way.

The result, as far as Alice is concerned, is that Bob's system finally resides predictably in the state $|\psi\rangle$ †.

How can Alice convince herself that such is the case? Well, if Bob is willing to reveal his location, she just need walk to his site and perform the **YES–NO** measurement: $|\psi\rangle\langle\psi|$ versus $I - |\psi\rangle\langle\psi|$. The outcome will be a **YES** with probability one for her if all has gone well in carrying out the protocol. Thus, for the cost of a measurement on a causally disconnected system and two bits worth of causal action on the system of actual interest—i.e. one of the four Pauli rotations—Alice can sharpen her predictability to complete certainty for *any* **YES–NO** observable she wishes.

Penrose argues in his book *The Emperor's New Mind* [27] that when a system 'has' a state $|\psi\rangle$ there ought to be some property in the system (in and of itself) that corresponds to its ' $|\psi\rangle$ 'ness'. For how else could the system be prepared to reveal a **YES** in the case that Alice actually checks it? Asking this rhetorical question with a sufficient amount of command is enough to make many a would-be informationist weak at the knees. But there is a crucial oversight implicit in its confidence, and we have caught it in action. If Alice fails to reveal her information to anyone, there is no one else in the world who can predict the qubit's ultimate revelation with certainty—examining Alice's measurement device before and after the measurement tells us nothing about the posterior quantum state Alice ends up with for Bob's system. How can a secret be anything more than pure information? More importantly, there is nothing in quantum mechanics that gives the qubit the power to stand up and say **YES** all by itself: if Alice does not take the time to walk over to it and interact with it, there is no revelation. There is only the confidence in Alice's mind that, *should* she interact with it, she *could* predict the consequence‡ of that interaction.

4. Information about what?

There are great rewards in being a new parent. Not least of all is the opportunity to have a close-up look at a mind in formation. Last year, I watched my two-year old learn things at a fantastic rate, and though there were untold lessons for her, there were a sprinkling for me too. For instance, I started to see her come to grips with the idea that there is a world independent of her desires. What struck me was the contrast between that and the gain of confidence I saw grow in her that there are aspects of existence she *could* control. The two go hand in hand. She pushes on the world, and sometimes it gives in a way that she has learned to predict, and sometimes it pushes back in a way she has not foreseen (and may never be able to). If she could manipulate the world to the complete desires of her will, there would be little difference between wake and dream.

The main point is that she learns from her forays into the world. In my cynical moments, I find myself thinking, 'How can she think that she's learned anything at all? She has no theory of measurement. She leaves measurement completely undefined. How can she have a stake to knowledge if she does not have a theory of how she learns?'

† As far as Bob is concerned, nothing whatsoever changes about the system in his possession: it started in the completely mixed state $\rho = \frac{1}{2}I$ and remains that way.

‡ I adopt this terminology to be similar to Savage's book [28], on rational decision theory.

Hideo Mabuchi once told me, ‘The quantum measurement problem refers to a set of people.’ And though that is a bit harsh, maybe it also contains a bit of the truth. With the physics community making use of theories that tend to last between 100 and 300 years, we are apt to forget that scientific views of the world are built from the top down, not from the bottom up. The experiment is the basis of all that we try to describe with science. But an experiment is an active intervention into the course of nature on the part of the experimenter; it is not contemplation of nature from afar. We set up this or that experiment to see how nature reacts. It is the conjunction of myriads of such interventions and their consequences that we record into our data books.

We tell ourselves that we have learned something new when we can distill from the data a compact description of all that was seen and—even more tellingly—when we can dream up further experiments to corroborate that description. This is the minimal requirement of science. If, however, from such a description we can *further* distill a model of a free-standing ‘reality’ independent of our interventions, then so much the better. I have no bone to pick with reality. It is the most solid thing we can hope for from a theory. Classical physics is the ultimate example in that regard. It gives us a compact description, but it can give much more if we want it to.

The thing to realize, however, is that there is no logical necessity that such a worldview should always be obtainable. If the world is such that we can never identify a reality—a *free-standing* reality—independent of our experimental interventions, then we must be prepared for that too. That is where quantum theory in its most minimal and conceptually simplest dispensation seems to stand [29]. It is a theory whose terms refer predominantly to our interface with the world. It is a theory that cannot go the extra step that classical physics did without tearing and straining to rhyme. It is a theory not about observables, not about beables, but about ‘ringables’. We tap a bell with our gentle touch and listen for its beautiful ring.

So what are the ways we can intervene on the world? What are the ways we can push it and wait for its unpredictable reaction? The textbook story is that measurables correspond to Hermitian operators. Or to say it in more modern language, to each observable there corresponds a set of orthogonal projection operators $\{\Pi_i\}$ over a complex Hilbert space \mathcal{H}_D which form a complete resolution of the identity, $\sum_i \Pi_i = I$. The index i labels the potential outcomes of the measurement (or *intervention*, to slip back into the language above). When an observer possesses the information ρ —captured most generally by a mixed-state density operator—quantum mechanics dictates that he can expect the various outcomes with a probability $P(i) = \text{tr}(\rho\Pi_i)$.

The best justification for this probability rule comes by way of Gleason’s amazing 1957 theorem [19]. For it states that the standard rule is the *only* rule that satisfies a very simple kind of non-contextuality for measurement outcomes [30]. In particular, if one contemplates measuring two distinct observables $\{\Pi_i\}$ and $\{\Gamma_i\}$ which happen to share a single projector Π_k , then the probability of outcome k is independent of which observable it is associated with. More formally, the statement is this. Let \mathcal{P}_D be the set of projectors associated with a (real or complex) Hilbert space \mathcal{H}_{rmd} for $D \geq 3$, and let $f : \mathcal{P}_D \rightarrow [0, 1]$ be such that $\sum_i f(\Pi_i) = 1$ whenever a set of projectors $\{\Pi_i\}$ forms an observable. The theorem concludes that there exists a density operator ρ such that $f(\Pi) = \text{tr}(\rho\Pi)$. In fact, in a single blow,

Gleason's theorem derives not only the probability rule, but also the state-space structure for quantum mechanical states (i.e. that it corresponds to the convex set of density operators).

In itself this is no small feat, but the thing that makes the theorem an 'amazing' theorem is the sheer difficulty required to prove it [31]. Note that no restrictions have been placed upon the function f beyond the ones mentioned above. There is no assumption that it need be differentiable, nor that it even need be continuous. All of that, and linearity too, comes from the structure of the observables—i.e. that they are complete sets of orthogonal projectors onto a linear vector space.

Nonetheless, one should ask: does this theorem really give *the physicist* a clearer vision of where the probability rule comes from? Astounding feats of mathematics are one thing; insight into physics is another. The two are often at opposite ends of the spectrum. As fortunes turn, a unifying strand can be drawn by viewing quantum foundations in the light of quantum information.

The place to start is to drop the fixation that the basic set of observables in quantum mechanics are complete sets of orthogonal projectors. In quantum information theory it has been found to be extremely convenient to expand the notion of measurement to also include general positive operator-valued measures (POVMs) [24, 32]. In other words, in place of the usual textbook notion of measurement, *any* set $\{E_d\}$ of positive-semi-definite operators on \mathcal{H}_D that forms a resolution of the identity—i.e. that satisfies $\langle \psi | E_d | \psi \rangle \geq 0$ for all $|\psi\rangle \in \mathcal{H}_D$ and $\sum_d E_d = I$ —counts as a measurement. The outcomes of the measurement are identified with the indices d , and the probabilities of the outcomes are computed according to a generalized Born rule,

$$P(d) = \text{tr}(\rho E_d). \quad (1)$$

The set $\{E_d\}$ is called a POVM, and the operators E_d are called POVM elements. (In the non-standard language promoted earlier, the set $\{E_d\}$ signifies an intervention into nature, while the individual E_d represent the potential consequences of that intervention.) Unlike standard measurements, there is no limitation on the number of values the index d can take. Moreover, the E_d may be of any rank, and there is no requirement that they be mutually orthogonal.

The way this expansion of the notion of measurement is *usually* justified is that any POVM can be represented formally as a standard measurement on an ancillary system that has interacted in the past with the system of actual interest. Indeed, suppose the system and ancilla are initially described by the density operators ρ_S and ρ_A respectively. The conjunction of the two systems is then described by the initial quantum state $\rho_{SA} = \rho_S \otimes \rho_A$. An interaction between the systems via some unitary time evolution leads to a new state $\rho_{SA} \rightarrow U \rho_{SA} U^\dagger$. Now, imagine a standard measurement on the ancilla. It is described on the total Hilbert space via a set of orthogonal projection operators $\{I \otimes \Pi_d\}$. An outcome d will be found with probability

$$P(d) = \text{tr}(U(\rho_S \otimes \rho_A)U^\dagger(I \otimes \Pi_d)). \quad (2)$$

The number of outcomes in this seemingly indirect notion of measurement is limited only by the dimensionality of the ancilla's Hilbert space—in principle, there can be arbitrarily many.

As advertised, it turns out that the probability formula above can be expressed in terms of operators on the system’s Hilbert space alone: this is the origin of the POVM. If we let $|s_\alpha\rangle$ and $|a_c\rangle$ be an orthonormal basis for the system and ancilla respectively, then $|s_\alpha\rangle|a_c\rangle$ will be a basis for the composite system. Using the cyclic property of the trace in equation (2), we obtain

$$P(d) = \sum_\alpha \langle s_\alpha | \rho_S \left(\sum_c \langle a_c | \left((I \otimes \rho_A) U^\dagger (I \otimes \Pi_d) U \right) | a_c \rangle \right) | s_\alpha \rangle. \tag{3}$$

Letting tr_A and tr_S denote partial traces over the system and ancilla, respectively, it follows that $P(d) = \text{tr}_S(\rho_S E_d)$, where

$$E_d = \text{tr}_A \left((I \otimes \rho_A) U (I \otimes \Pi_d) U^\dagger \right) \tag{4}$$

is an operator acting on the Hilbert space of the original system. This proves half of what is needed, but it is also straightforward to go in the reverse direction—i.e. to show that for any POVM $\{E_d\}$, one can pick an ancilla and find operators ρ_A , U and Π_d such that equation (4) is true.

Putting this all together, there is a sense in which standard measurements capture everything that can be said about quantum measurement theory. What I would like to bring up is whether this standard way of justifying the POVM is the most productive point of view one can take. Might any of the mysteries of quantum mechanics be alleviated by taking the POVM as a basic notion of measurement? Does the POVM’s utility portend a larger role for it in the foundations of quantum mechanics?

Standard measurements	Generalized measurements
$\{\Pi_i\}$	$\{E_d\}$
$\langle \psi \Pi_i \psi \rangle \geq 0, \forall \psi\rangle$	$\langle \psi E_d \psi \rangle \geq 0, \forall \psi\rangle$
$\sum_i \Pi_i = I$	$\sum_d E_d = I$
$P(i) = \text{tr}(\rho \Pi_i)$	$P(d) = \text{tr}(\rho E_d)$
$\Pi_i \Pi_j = \delta_{ij} \Pi_i$	—

I try to make this point dramatic in my lectures by exhibiting the table above. On the left-hand side there is a list of various properties for the standard notion of a quantum measurement. On the right-hand side, there is an almost identical list of properties for the POVMs. The only difference between the two columns is that the right-hand one is *missing* the orthonormality condition required of a standard measurement. The question I ask the audience is this: does the addition of that one extra assumption really make the process of measurement any less mysterious? Indeed, I imagine myself teaching quantum mechanics for the first time and taking a vote with the best audience of all, the students. ‘Which set of postulates for quantum measurement would you prefer?’ I am quite sure they would respond with a blank stare. But that is the point! It would make no difference to them, and it should make no difference to us. The only issue worth debating is which notion of measurement will allow us to see more deeply into quantum mechanics.

Therefore let us pose the question Gleason did, but with POVMs. In other words, let us suppose that the ways an experimenter can intervene on a quantum system corresponds to the set of POVMs on its Hilbert space \mathcal{H}_D . It is the task of the theory to give him probabilities for the consequences of his interventions. Concerning those probabilities, let us again assume non-contextuality—i.e. whatever the probability for a given consequence E_c is, it does not depend upon whether E_c is associated with the POVM $\{E_d\}$ or, instead, any other one $\{\tilde{E}_d\}$. This means there exists a function $f: \mathcal{E}_D \rightarrow [0, 1]$, where $\mathcal{E}_D = \{E: 0 \leq \langle \psi|E|\psi \rangle \leq 1, \forall |\psi\rangle \in \mathcal{H}_D\}$, such that whenever $\{E_d\}$ forms a POVM, $\sum_d f(E_d) = 1$. (In general, we will call any function satisfying $f(E) \geq 0$ and $\sum_d f(E_d) = \text{constant}$ a *frame function*, in analogy to Gleason's non-negative frame functions.)

It will come as no surprise, of course, that a Gleason-like theorem must hold for the function f . Namely, it can be shown that there must exist a density operator ρ for which $f(E) = \text{tr}(\rho E)$. This was recently shown by Busch [17] and, independently, by Renes and collaborators [18]. What is surprising however is the utter simplicity of the proof.

To show that off, let us exhibit the whole proof for the special case where \mathcal{H}_D is defined over the field of (complex) rational numbers. The full theorem based on the continuum, is a minor extension of this. It is no problem to see that f is 'linear' with respect to positive combinations of operators that never go outside \mathcal{E}_D . Consider a three-element POVM $\{E_1, E_2, E_3\}$. By assumption $f(E_1) + f(E_2) + f(E_3) = 1$. However, we can also group the first two elements in this POVM to obtain a new POVM, and must therefore have $f(E_1 + E_2) + f(E_3) = 1$. In other words, the function f must be additive with respect to a fine-graining operation:

$$f(E_1 + E_2) = f(E_1) + f(E_2). \quad (5)$$

Similarly for any two integers m and n , $f(E) = mf[(1/m)E] = nf[(1/n)E]$. Suppose $n/m \leq 1$. Then if we write $E = nG$, this statement becomes: $f[(n/m)G] = (n/m)f(G)$. In other words, we immediately have a kind of limited linearity on \mathcal{E}_D .

One *might* imagine using this property to cap off the theorem in the following way. Clearly the full D^2 -dimensional vector space \mathcal{O}_D of Hermitian operators on \mathcal{H}_D is spanned by the set \mathcal{E}_D since that set contains, among other things, all the projection operators. Thus, we can write any operator $E \in \mathcal{E}_D$ as a linear combination $E = \sum_i \alpha_i E_i$ for some fixed operator-basis $\{E_i\}_{i=1}^{D^2}$. 'Linearity' of f would then give $f(E) = \sum_i \alpha_i f(E_i)$. So, if we define ρ by solving the D^2 linear equations $\text{tr}(\rho E_i) = f(E_i)$, we would have

$$f(E) = \sum_i \alpha_i \text{tr}(\rho E_i) = \text{tr}\left(\rho \sum_i \alpha_i E_i\right) = \text{tr}(\rho E) \quad (6)$$

and essentially be done. (Positivity and normalization of f would require ρ to be an actual density operator.) The *problem* is that in the expansion of E there is no guarantee that the coefficients α_i can be chosen so that $\alpha_i E_i \in \mathcal{E}_D$.

What remains to be shown is that f can be extended uniquely to a function that is truly linear on \mathcal{O}_D . This too is rather simple. First, take any positive semi-definite operator E . We can always find a positive rational number g such that

$E = gG$ and $G \in \mathcal{E}_D$. Therefore, we can simply define $f(E) \equiv gf(G)$. To see that this definition is unique, suppose there are two such operators G_1 and G_2 (with corresponding numbers g_1 and g_2) such that $E = g_1G_1 = g_2G_2$. Further suppose $g_2 \geq g_1$. Then $G_2 = (g_1/g_2)G_1$ and, by the homogeneity of the original unextended definition of f , we obtain $g_2f(G_2) = g_1f(G_1)$. Furthermore this extension retains the additivity of the original function. For suppose that neither E nor G , although positive semi-definite, are necessarily in \mathcal{E}_D . We can find a positive rational number $c \geq 1$ such that $(1/c)(E + G)$, $(1/c)E$, and $(1/c)G$ are all in \mathcal{E}_D . Then, by the rules we have already obtained,

$$f(E + G) = cf\left(\frac{1}{c}(E + G)\right) = cf\left(\frac{1}{c}E\right) + cf\left(\frac{1}{c}G\right) = f(E) + f(G). \quad (7)$$

Let us now further extend f 's domain to the full space \mathcal{O}_d . This can be done by noting that any operator H can be written as the difference $H = E - G$ of two positive semi-definite operators. Therefore define $f(H) \equiv f(E) - f(G)$, from which it also follows that $f(-G) = -f(G)$. To see that this definition is unique suppose there are four operators E_1, E_2, G_1 and G_2 , such that $H = E_1 - G_1 = E_2 - G_2$. It follows that $E_1 + G_2 = E_2 + G_1$. Applying f (as extended in the previous paragraph) to this equation, we obtain $f(E_1) + f(G_2) = f(E_2) + f(G_1)$ so that $f(E_1) - f(G_1) = f(E_2) - f(G_2)$. Finally, with this new extension, full linearity can be checked immediately. This completes the proof as far as the (complex) rational number field is concerned: because f extends uniquely to a linear functional on \mathcal{O}_D , we can indeed go through the steps of equation (6) without worry.

There are two things that are significant about this proof. First, in contrast to Gleason's original theorem, there is nothing to bar the same logic from working when $D = 2$. This is quite nice because much of the community has gotten into the habit of thinking that there is nothing particularly 'quantum mechanical' about a single qubit. Indeed, because orthogonal projectors on \mathcal{H}_2 can be mapped onto antipodes of the Bloch sphere, it is known that the measurement-outcome statistics for any standard measurement can be mocked-up through a non-contextual hidden-variable theory. What this result shows is that this simply is not the case when one considers the full set of POVMs as one's potential measurements.

The other important thing is that the theorem works for Hilbert spaces over the rational number field: one does not need to invoke the full power of the continuum. This contrasts with the surprising result of Meyer [33] that the standard Gleason theorem fails in such a setting. The present theorem hints at a kind of resiliency to the structure of quantum mechanics which falls through the mesh of the standard Gleason result: the probability rule for POVMs does not actually depend so much upon the detailed workings of the number field.

Of course we are not getting something for nothing. The reason the present derivation is so easy in contrast to the standard proof is that *mathematically* the assumption of POVMs as the basic notion of measurement is significantly stronger than the usual assumption. *Physically*, though, I would say it is just the opposite. Why add extra restrictions to the notion of measurement when they only make the route from basic assumption to practical usage more circuitous than need be? In the end, a measurement, at some point in the chain, is always a black box—for the physicist every bit as much as for my two-year old daughter. The issue is only to choose the prettiest black box possible.

4.1. The International Bureau of Weights and Measures

There is one further, particularly important, advantage to thinking of POVMs as the basic notion of measurement in quantum mechanics. For with an appropriately chosen *single* POVM one can stop thinking of the quantum state as a linear operator altogether, and instead start thinking of it as a probabilistic judgement with respect to the (potential) outcomes of a *standard quantum measurement*. That is, a measurement device right next to the standard kilogram and the standard metre in a carefully guarded vault, deep within the bowels of the International Bureau of Weights and Measures†. Here is what I mean by this.

Our problem hinges on finding a measurement for which the probabilities of outcomes completely specify a unique density operator. Such measurements are called *informationally complete* and have been studied for some time [35]. Here however, the picture is most pleasing if we consider a slightly refined version of the notion—that of the *minimal* informationally complete measurement [36]. The space of Hermitian operators on \mathcal{H}_D is itself a linear vector space of dimension D^2 . The quantity $\text{tr}(A^\dagger B)$ serves as an inner product on that space. Hence, if we can find a POVM $\mathcal{E} = \{E_d\}$ consisting of D^2 linearly independent operators, the probabilities $P(d) = \text{tr}(\rho E_d)$ —now thought of as *projections* in the directions of the vectors E_d —will completely specify the operator ρ . Any two distinct density operators ρ and σ must give rise to distinct outcome statistics for this measurement. The minimal number of outcomes a POVM can have and still be informationally complete is D^2 .

Do minimal informationally complete POVMs exist? The answer is yes. Here is a simple way to produce one, although there are many other ways. Start with a complete orthonormal basis $|e_j\rangle$ on \mathcal{H}_D . One can check that the following D^2 rank-1 projectors Π_d form a linearly independent set.

- (1) For $d = 1, \dots, D$, let $\Pi_d = |e_j\rangle\langle e_j|$, where j , too, runs over the values $1, \dots, D$.
- (2) For $d = D + 1, \dots, \frac{1}{2}D(D + 1)$, let $\Pi_d = \frac{1}{2}(|e_j\rangle + |e_k\rangle)(\langle e_j| + \langle e_k|)$, where $j < k$.
- (3) For $d = \frac{1}{2}D(D + 1) + 1, \dots, D^2$, let $\Pi_d = \frac{1}{2}(|e_j\rangle + i|e_k\rangle)(\langle e_j| - i\langle e_k|)$, where again $j < k$.

All that remains is to transform these linearly independent operators Π_d into a proper POVM. This can be done by considering a positive semi-definite operator G defined by

$$G = \sum_{d=1}^{D^2} \Pi_d. \quad (8)$$

It is straightforward to show that $\langle \psi | G | \psi \rangle > 0$ for all $|\psi\rangle \neq 0$, thus establishing that G is positive definite and hence invertible. Applying the (invertible) linear transformation $X \rightarrow G^{-1/2} X G^{-1/2}$ to equation (8), we find a valid decomposition of the identity,

$$I = \sum_{d=1}^{D^2} G^{-1/2} \Pi_d G^{-1/2}. \quad (9)$$

† This idea has its roots in L. Hardy's important paper [34].

The operators $E_d = G^{-1/2} \Pi_d G^{-1/2}$ satisfy the conditions of a POVM, and moreover, they retain the rank and linear independence of the original Π_d . Thus we have what we need.

With the existence of minimal informationally complete POVMs assured, we can think about the vault in Paris. Let us suppose from here out that it contains a machine which enacts a minimal informationally complete POVM E_h whenever it is used. We reserve the index h to denote the outcomes of this standard quantum measurement, for they will replace the notion of the ‘hypothesis’ in classical statistical theory. Let us develop this from a Bayesian point of view.

Whenever one has a quantum system in mind, it is legitimate for him to use all he knows and believes of it to ascribe a probability function $P(h)$ to the (potential) outcomes of this standard measurement. In fact, that is all a quantum state *is* from this point of view: it is a subjective judgement about which consequence will obtain as a result of an interaction between *one’s* system and *that* machine. Whenever one performs a measurement $\{E_d\}$ on the system—one different from the standard quantum measurement $\{E_h\}$ —at the most basic level of understanding, all one is doing is gathering (or evoking) a piece of data d that (among other things) allows one to update from one’s initial subjective judgement $P(h)$ to *something* else $P_d(h)$.

It is important to recognize that, with this change of description, we may be edging toward a piece of quantum mechanics that is not of information theoretic origin. It is this. If one accepts quantum mechanics and supposes that one has a system for which the standard quantum measurement device has D^2 outcomes, then one is no longer free to make just any subjective judgement $P(h)$ he pleases. There are constraints. Let us call the allowed region of initial judgements $\mathcal{P}_{\text{SQM}} \dagger$.

For any minimal informationally complete POVM $\{E_h\}$, $P(h)$ must be bounded away from unity for all its possible outcomes. Thus even at this stage, there is something driving a wedge between quantum mechanics and simple Bayesian probability theory. When one accepts quantum mechanics, one voluntarily accepts a restriction on one’s subjective judgements for the consequences of a standard quantum measurement intervention: for all consequences h , there are no conditions whatsoever convincing enough to compel *one* to a probability ascription $P(h) = 1$. That is, one gives up on the hope of certainty. This, indeed, one might pinpoint as an assumption about the physical world that goes beyond pure probability theory ‡.

But what is that assumption in physical terms? What is our best description of the wedge? Some think they already know the answer, and it is quantum entanglement.

5. Wither entanglement?

Quantum entanglement has certainly captured the attention of our community. By most accounts it is the main ingredient in quantum information theory and

† This region is a convex set. For details about the convex region generated by an arbitrary POVM, see [37].

‡ It is at this point that the present account of quantum mechanics differs most crucially from [34]. Hardy sees quantum mechanics as a generalization and extension of classical probability theory, whereas quantum mechanics is depicted here as a restriction to probability theory. It is a restriction that takes into account how we ought to think and gamble in light of a certain physical fact—a fact which we are working to identify.




Figure 1. The planar surface represents the space of all probability distributions over D^2 outcomes. Accepting quantum mechanics is, in part, accepting that one's subjective beliefs for the outcomes of a standard quantum measurement device will not fall outside a certain convex set. Each point within the convex region represents a valid quantum state.

quantum computing, and it is the main mystery of the quantum foundations. But what is it? Where does it come from?

The predominant purpose it has served in this paper has been as a kind of background. For it, more than any other ingredient in quantum mechanics, has clinched the issue of 'information about what?' in the author's mind: that information cannot be about a pre-existing reality (a hidden variable) unless we are willing to renege on our reason for rejecting the quantum state's objective reality in the first place. What I am alluding to here is the conjunction of the Einstein argument reported in section 3 and the phenomena of the Bell inequality violations by quantum mechanics. Putting those points together gave us that the information symbolized by a $|\psi\rangle$ must be information about the potential consequences of our interventions into the world.

But, now I would like to turn the tables and ask whether the structure of our potential interventions—the POVMs—can tell us something about the origin of entanglement. Could it be that the concept of entanglement is just a minor addition to the much deeper point that measurements have this structure?

The technical translation of this question is, why do we combine systems according to the tensor-product rule? There are certainly innumerable ways to combine two Hilbert spaces \mathcal{H}_A and \mathcal{H}_B to obtain a third \mathcal{H}_{AB} . We could take the direct sum of the two spaces $\mathcal{H}_{AB} = \mathcal{H}_A \oplus \mathcal{H}_B$. We could take their Grassmann product $\mathcal{H}_{AB} = \mathcal{H}_A \wedge \mathcal{H}_B$. We could take scads of other things. But instead we take their tensor product, $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. Why?

Could it arise from the selfsame considerations as of the previous section—namely, from a non-contextuality property for measurement-outcome probabilities? The answer is yes, and the theorem I am about demonstrate owes much in inspiration to [38].

Here is the scenario. Suppose we have two quantum systems, and we can make a measurement on each. On the first, we can measure any POVM on the D_A -dimensional Hilbert space \mathcal{H}_A ; on the second, we can measure any POVM on the D_B -dimensional Hilbert space \mathcal{H}_B . (This, one might think, is the very essence of having *two* systems rather than one—i.e. that we can probe them independently.) Moreover, suppose we may condition the second measurement on the nature and the outcome of the first, and vice versa. That is to say—walking from A to B —we could first measure $\{E_i\}$ on A , and then, depending on the outcome i , measure $\{F_j^i\}$ on B . Similarly—walking from B to A —we could first measure $\{F_j\}$ on B , and

then, depending on the outcome j , measure $\{E_i^j\}$ on A . So that we have valid POVMs, we must have $\sum_i E_i = I$ and $\forall i \sum_j F_j^i = I$, and similarly $\forall j \sum_i E_i^j = I$ and $\sum_j F_j = I$. Let us denote by S_{ij} an ordered pair of operators, either of the form (E_i, F_j^i) or of the form (E_i^j, F_j) , as appearing above. Let us call a set of such operators $\{S_{ij}\}$ a locally-measurable POVM tree.

Suppose now that—just as with the POVM-version of Gleason’s theorem in section 4—the joint probability $P(i, j)$ for the outcomes of such a measurement should not depend upon which tree S_{ij} is embedded in: this is essentially the same assumption we made there, but now applied to local measurements on the separate systems. In other words, let us suppose there exists a function $f : \mathcal{E}_{D_A} \times \mathcal{E}_{D_B} \rightarrow [0, 1]$ such that

$$\sum_{ij} f(S_{ij}) = 1 \tag{10}$$

whenever the S_{ij} form a locally-measurable POVM tree. Note in particular that this definition makes no use of the tensor product: the domain of f is the Cartesian product of \mathcal{E}_{D_A} and \mathcal{E}_{D_B} .

The theorem is this: if f satisfies equation (10) for all locally-measurable POVM trees, then there exists a linear operator \mathcal{L} on $\mathcal{H}_A \otimes \mathcal{H}_B$ such that

$$f(E, F) = \text{tr}(\mathcal{L}(E \otimes F)). \tag{11}$$

If \mathcal{H}_A and \mathcal{H}_B are defined over the field of complex numbers, then \mathcal{L} is unique. Uniqueness does not hold, however, if the underlying field is the real numbers.

The proof of this statement is almost a trivial extension of the proof in section 4. One again starts by showing additivity, but this time in the two variables E and F separately. For instance, for a fixed $E \in \mathcal{E}_{D_A}$, define $g_E(F) = f(E, F)$, and consider two locally-measurable POVM trees $\{(I - E, F_i), (E, G_\alpha)\}$ and $\{(I - E, F_i), (E, H_\beta)\}$, where $\{F_i\}$, $\{G_\alpha\}$ and $\{H_\beta\}$ are arbitrary POVMs on \mathcal{H}_B . Then equation (10) requires that $\sum_i g_{I-E}(F_i) + \sum_\alpha g_E(G_\alpha) = 1$ and $\sum_i g_{I-E}(F_i) + \sum_\beta g_E(H_\beta) = 1$. From this it follows that, $\sum_\alpha g_E(G_\alpha) = \sum_\beta g_E(H_\beta) = \text{constant}$. In other words, $g_E(F)$ is a frame function in the sense of section 4. Consequently, we know that we can use the same methods as there to uniquely extend $g_E(F)$ to a linear functional on the complete set of Hermitian operators on \mathcal{H}_B . Similarly, for fixed $F \in \mathcal{E}_{D_B}$, we can define $h_F(E) = f(E, F)$, and prove that this function too can be extended uniquely to a linear functional on the Hermitian operators on \mathcal{H}_A .

The linear extensions of $g_E(F)$ and $h_F(E)$ can be put together in a simple way to give a full bilinear extension to the function $f(E, F)$. Namely, for any two Hermitian operators A and B on \mathcal{H}_A and \mathcal{H}_B , respectively, let $A = \alpha_1 E_1 - \alpha_2 E_2$ and $B = \beta_1 F_1 - \beta_2 F_2$ be decompositions such that $\alpha_1, \alpha_2, \beta_1, \beta_2 \geq 0$, $E_1, E_2 \in \mathcal{E}_{D_A}$, and $F_1, F_2 \in \mathcal{E}_{D_B}$. Then define $f(A, B) \equiv \alpha_1 g_{E_1}(B) - \alpha_2 g_{E_2}(B)$. To see that this definition is unique, take any other decomposition $A = \tilde{\alpha}_1 \tilde{E}_1 - \tilde{\alpha}_2 \tilde{E}_2$. Then it is a simple matter to check that $\tilde{\alpha}_1 g_{\tilde{E}_1}(B) - \tilde{\alpha}_2 g_{\tilde{E}_2}(B) = \alpha_1 g_{E_1}(B) - \alpha_2 g_{E_2}(B)$ as desired.

With bilinearity for the function f established, we have essentially the full story. For, let $\{E_i\}$, $i = 1, \dots, D_A^2$, be a complete basis for the Hermitian operators on \mathcal{H}_A and let $\{F_j\}$, $j = 1, \dots, D_B^2$, be a complete basis for the Hermitian operators

on \mathcal{H}_B . If $E = \sum_i \alpha_i E_i$ and $F = \sum_j \beta_j F_j$, then $f(E, F) = \sum_{ij} \alpha_i \beta_j f(E_i, F_j)$. Define \mathcal{L} to be a linear operator on $\mathcal{H}_A \otimes \mathcal{H}_B$ satisfying the $(D_A D_B)^2$ linear equations

$$\text{tr} \left(\mathcal{L}(E_i \otimes F_j) \right) = f(E_i, F_j). \quad (12)$$

Such an operator always exists. Consequently we have,

$$f(E, F) = \sum_{ij} \alpha_i \beta_j \text{tr} \left(\mathcal{L}(E_i \otimes F_j) \right) = \text{tr} \left(\mathcal{L}(E \otimes F) \right). \quad (13)$$

For complex Hilbert spaces \mathcal{H}_A and \mathcal{H}_B , the uniqueness of \mathcal{L} follows because the set $\{E_i \otimes F_j\}$ forms a complete basis for the Hermitian operators on $\mathcal{H}_A \otimes \mathcal{H}_B$ [39]. For real Hilbert spaces, however, the analogue of the Hermitian operators are the symmetric operators. The dimensionality of the space of symmetric operators on a real Hilbert space \mathcal{H}_D is $\frac{1}{2}D(D+1)$, rather than the D^2 it is for the complex case. This means that in the steps above only $\frac{1}{4}D_A D_B (D_A + 1)(D_B + 1)$ equations will appear in equation (12), whereas $\frac{1}{2}D_A D_B (D_A D_B + 1)$ are needed to uniquely specify \mathcal{L} .

This establishes the theorem. It would be nice if we could go further and establish the full probability rule for local quantum measurements—i.e. that \mathcal{L} must be a density operator. Unfortunately, our assumptions are not strong enough for that. For instance, suppose $\mathcal{H}_A = \mathcal{H}_B$ and let Φ be any positive, but not completely positive, trace-preserving linear map from and to the operators on \mathcal{H}_A . Then, for any maximally entangled state $|\psi_{ME}\rangle$ on $\mathcal{H}_A \otimes \mathcal{H}_B$, $\mathcal{L} = I \otimes \Phi(|\psi_{ME}\rangle\langle\psi_{ME}|)$ will not be a density operator. Yet, the quantity $f(E, F)$ in equation (11) will nevertheless be non-negative [40].

Of course, one could recover positivity for \mathcal{L} by requiring that it give positive probabilities even for non-local measurements (i.e. resolutions of the identity operator on $\mathcal{H}_A \otimes \mathcal{H}_B$). But in the purely local setting contemplated here, that would be a cheap way out. For, one should ask in good conscience what ought to be the rule for defining the full class of measurements (including non-local measurements): why should it correspond to an arbitrary resolution of the identity on the tensor product? There is nothing that makes it obviously so, unless one has already accepted standard quantum mechanics. Alternatively, it must be possible to give a purely local condition that will restrict \mathcal{L} to be a density operator. This is because \mathcal{L} , as noted above, is uniquely determined by the function $f(E, F)$; we never need to look further than the probabilities of local measurements outcomes in specifying \mathcal{L} . Ferreting out such a condition supplies an avenue for future research.

All of this does not, however, take away from the fact that whatever \mathcal{L} is, it must be a linear operator on the tensor product $\mathcal{H}_A \otimes \mathcal{H}_B$. Therefore, let us close by emphasizing the striking feature of this way of deriving the tensor-product rule for combining separate quantum systems: it is built on the very concept of local measurement. There is nothing ‘spooky’ or ‘non-local’ about it.

Thus, the tensor-product rule, and with it quantum entanglement, seems to be more a statement of locality than anything else. It, like the single-system probability rule, is more a product of the structure of the observables—that they are POVMs—combined with non-contextuality. In searching for the secret ingredient to drive a

wedge between general Bayesian probability theory and quantum mechanics, it seems that the direction *not* to look is toward quantum entanglement. Perhaps the trick instead is to dig deeper into the Bayesian toolbox.

6. Whither Bayes' rule?

Quantum states are states of information, knowledge, belief, pragmatic gambling commitments, *not* states of nature. That statement is the cornerstone of this paper. Thus, in searching to make sense of the remainder of quantum mechanics, one strategy ought to be to seek guidance from the most developed avenue of 'rational-decision theory' to date—Bayesian probability theory [41]. Indeed, the very aim of Bayesian theory is to develop reliable methods of reasoning and making decisions in the light of incomplete information. To what extent does that structure mesh with the seemingly independent structure of quantum mechanics?

The core of the matter is the manner in which states of belief are updated in the two theories. At first sight, they appear to be quite different in character. To see this, let us first explore how quantum mechanical states change when information is gathered. In older accounts of quantum mechanics, one often encounters the 'collapse postulate' as a basic statement of the theory. One hears things like, 'Axiom 5: upon the completion of an ideal measurement of an Hermitian operator H , the system is left in an eigenstate of H .' In quantum information, however, it has become clear that it is useful to broaden the notion of measurement, and with it, the analysis of how a state can change in the process. The foremost reason for this is that the collapse postulate is simply not true in general: depending upon the exact nature of the measurement interaction, there may be any of a large set of possibilities for the final state of a system.

The broadest notion of state change is this [32]. Suppose one's initial state for a quantum system is a density operator ρ , and a POVM $\{E_d\}$ is measured on that system. Then, the state after the measurement can be *any* state ρ_d of the form

$$\rho_d = \frac{1}{\text{tr}(\rho E_d)} \sum_i A_{di} \rho A_{di}^\dagger, \quad (14)$$

where

$$\sum_i A_{di}^\dagger A_{di} = E_d. \quad (15)$$

Note the immense generality of this formula. There is no constraint on the number of indices i in the A_{di} and these operators need not even be Hermitian.

The *usual* justification for this kind of generality—just as in the case of the common justification for the POVM formalism—comes by imagining that the measurement arises in an indirect fashion, rather than as a direct observation. In other words, the primary system is pictured to interact with an ancilla first, and only then subjected to a 'real' measurement on the ancilla alone. The trick is that one posits a kind of projection postulate on the primary system due to this process. This assumption has a much safer feel than the raw projection postulate since, after the interaction, no measurement on the ancilla should cause a physical perturbation to the primary system.

More formally, we can start out by following the usual derivation of how a POVM can be thought of as a standard measurement on a larger system, but in place of equation (2) we must make an assumption about how the system's state changes. For this one invokes a kind of 'projection-postulate-at-a-distance'. Namely, one takes

$$\rho_d = \frac{1}{P(d)} \text{tr}_A ((I \otimes \Pi_d)U(\rho_S \otimes \rho_A)U^\dagger(I \otimes \Pi_d)). \tag{16}$$

The reason for invoking the partial trace is to make sure that any hint of a state change for the ancilla remains unaddressed.

To see how expression (16) makes connection to equation (14), denote the eigenvalues and eigenvectors of ρ_A by λ_α and $|a_\alpha\rangle$ respectively. Expanding equation (16), we have

$$\rho_d = \frac{1}{P(d)} \sum_{\alpha\beta} (\lambda_\alpha^{1/2} \langle a_\beta | (I \otimes \Pi_d) U^\dagger | a_\alpha \rangle) \rho_S (\langle a_\alpha | U (I \otimes \Pi_d) | a_\beta \rangle \lambda_\alpha^{1/2}). \tag{17}$$

A representation of the form in equation (14) can be made by taking $A_{b\alpha\beta} = \lambda_\alpha^{1/2} \langle a_\alpha | U (I \otimes \Pi_d) | a_\beta \rangle$ and lumping the two indices α and β into the single index i . Indeed, one can easily check that equation (15) holds. This completes what we had set out to show. However, just as with the case of the POVM $\{E_d\}$, one can always find a way to reverse engineer the derivation: given a set of A_{di} , one can always find a U , a ρ_A , and set of Π_d such that equation (16) becomes true.

Of course the old collapse postulate is contained within the extended formalism as a special case: there, one just takes both sets $\{E_d\}$ and $\{A_{di} = E_d\}$ to be sets of orthogonal projection operators. Let us take a moment to think about this special case in isolation. What is distinctive about it is that it captures in the extreme a common folklore associated with the measurement process. For it tends to convey the image that measurement is a kind of gut-wrenching violence: in one moment the state is $\rho = |\psi\rangle\langle\psi|$, while in the very next it is a $\Pi_i = |i\rangle\langle i|$. Moreover, such a wild transition need depend upon no details of $|\psi\rangle$ and $|i\rangle$; in particular the two states may even be almost orthogonal to each other. In density-operator language, there is no sense in which Π_i is contained in ρ : the two states are in distinct places of the operator space. That is, $\rho \neq \sum_i P(i)\Pi_i$.

Contrast this with the description of information gathering that arises in Bayesian probability theory. There, an initial state of belief is captured by a probability distribution $P(h)$ for some hypothesis H . The way gathering a piece of data d is taken into account in assigning one's new state of belief is through Bayes' conditionalization rule. That is to say, one expands $P(h)$ in terms of the relevant joint probability distribution and picks off the appropriate term:

$$P(h) = \sum_d P(d)P(h|d) \tag{18}$$

$$P(h) \xrightarrow{d} P(h|d). \tag{19}$$

How gentle this looks in comparison to quantum collapse! When one gathers new information, one simply refines one's old beliefs in the most literal of senses. It is

not as if the new state is incommensurable with the old. *It was always there*; it was just initially averaged in with various other potential beliefs.

Why does quantum collapse not look more like Bayes' rule? Is quantum collapse really a more violent kind of change, or might it be an artefact of a problematic representation? By this stage, it should come as no surprise to the reader that dropping the ancilla from our image of generalized measurements will be the first step to progress. Taking the transition from ρ to ρ_d in equations (14) and (15) as the basic statement of what quantum measurement *is* is a good starting point.

To accentuate a similarity between equation (14) and Bayes' rule, let us first contemplate cases of it where the index i takes on a single value. Then, we can conveniently drop that index and write

$$\rho_d = \frac{1}{P(d)} A_d \rho A_d^\dagger, \tag{20}$$

where

$$E_d = A_d^\dagger A_d. \tag{21}$$

In a loose way, one can say that measurements of this sort are the most efficient they can be for a given POVM $\{E_d\}$: for, a measurement interaction with an explicit i -dependence may be viewed as 'more truly' a measurement of a finer-grained POVM that just happens to throw away some of the information it gained. Let us make this point more precise.

Notice that Bayes' rule has the property that one's uncertainty about a hypothesis can be expected to decrease upon the acquisition of data. This can be made rigorous, for instance, by gauging uncertainty in terms of the Shannon entropy function, $S(H) = -\sum_h P(h) \log P(h)$. Since $f(x) = -x \log x$ is concave on the interval $[0, 1]$, it follows that

$$S(H) \geq -\sum_d P(d) \sum_h P(h|d) \log P(h|d) = \sum_d P(d) S(H|d). \tag{22}$$

Indeed we hope to find a similar statement for how the result of efficient quantum measurements decrease uncertainty. But, what can be meant by a decrease of uncertainty through quantum measurement? I have argued that the information gain in a measurement cannot be about a pre-existing reality. The way out of the impasse is simple: the uncertainty that decreases in quantum measurement is the uncertainty one expects for the results of other potential measurements.

A good way to quantify this has to do with the von Neumann entropy, $S(\rho) = -\text{tr } \rho \log \rho$. The intuitive meaning of the von Neumann entropy can be found by first thinking about the Shannon entropy. Consider any von Neumann measurement \mathcal{P} consisting of d one-dimensional orthogonal projectors Π_i . A natural question to ask is: with respect to a given density operator ρ , which measurement \mathcal{P} will give the most predictability over its outcomes? As it turns out, the answer is any \mathcal{P} that forms a set of eigenprojectors for ρ [42]. When this is obtained, the Shannon entropy of the measurement outcomes reduces to simply the von Neumann entropy of the density operator. The von Neumann entropy, then, signifies the amount of unpredictability one achieves by way of a standard

measurement in a best case scenario. Indeed, true to one’s intuition, one has the most predictability by this account when ρ is a pure state—for then $S(\rho) = 0$. Alternatively, one has the least knowledge when ρ is proportional to the identity operator—for then any measurement \mathcal{P} will have outcomes that are all equally likely.

The way to get at a quantum statement of equation (22) is to make use of the fact that $S(\rho)$ is concave in the variable ρ [43]. A function F is concave in ρ when

$$F(t\tilde{\rho}_0 + (1 - t)\tilde{\rho}_1) \geq tF(\tilde{\rho}_0) + (1 - t)F(\tilde{\rho}_1), \tag{23}$$

for any density operators $\tilde{\rho}_0$ and $\tilde{\rho}_1$ and any real number $t \in [0, 1]$. Therefore, one might hope that

$$F(\rho) \geq \sum_d P(d)F(\rho_d). \tag{24}$$

Such a result, however, cannot arise in the easy fashion it did for the classical case of equation (22). This is because generally (as already emphasized), $\rho \neq \sum_d P(d)\rho_d$ for ρ_d defined as in equation (20). One has to be a little more roundabout to make a proof happen [43, 44].

For the purposes here, the key [43] is in noticing that

$$\rho = \rho^{1/2} I \rho^{1/2} = \sum_d \rho^{1/2} E_d \rho^{1/2} = \sum_d P(d)\tilde{\rho}_d, \tag{25}$$

where

$$\tilde{\rho}_d = \frac{1}{P(d)} \rho^{1/2} E_d \rho^{1/2} = \frac{1}{P(d)} \rho^{1/2} A_d^\dagger A_d \rho^{1/2}. \tag{26}$$

What is special about this decomposition of ρ is that for each d , ρ_d and $\tilde{\rho}_d$ have the same eigenvalues. This follows since $X^\dagger X$ and XX^\dagger have the same eigenvalues, for any operator X . In the present case, setting $X = A_d \rho^{1/2}$ does the trick. Using the fact that $S(\rho)$ depends only upon the eigenvalues of ρ we obtain:

$$S(\rho) \geq \sum_d P(d)S(\rho_d), \tag{27}$$

as we had been hoping for. Thus, in performing an efficient quantum measurement of a POVM $\{E_d\}$, an observer can expect to be left with less uncertainty than he started with.

In this sense, quantum ‘collapse’ does indeed have some of the flavour of Bayes’ rule. But we can expect more, and the derivation above hints at just the right ingredient: ρ_d and $\tilde{\rho}_d$ have the same eigenvalues! To see the impact of this, let us once again explore the content of equations (20) and (21). A common way to describe their meaning is to use the operator polar-decomposition theorem to rewrite equation (20) in the form

$$\rho_d = \frac{1}{P(d)} U_d E_d^{1/2} \rho E_d^{1/2} U_d^\dagger, \tag{28}$$

where U_d is a unitary operator. Since—subject only to the constraint of efficiency—the operators A_d are not determined any further than equation (21), U_d can be *any* unitary operator whatsoever. Thus, a customary way of thinking of the state-change process is to break it up into two conceptual pieces. First there is a ‘raw collapse’:

$$\rho \rightarrow \sigma_d = \frac{1}{P(d)} E_d^{1/2} \rho E_d^{1/2}. \tag{29}$$

Then, subject to the precise measurement interaction and the particular outcome d , one imagines the measuring device enforcing a further ‘back-action’ or ‘feedback’ on the measured system [45]:

$$\sigma_d \rightarrow \rho_d = U_d \sigma_d U_d^\dagger. \tag{30}$$

But this breakdown of the transition is a purely conceptual game.

Since the U_d are arbitrary to begin with, we might as well break down the state-change process into the following (non-standard) conceptual components. First one imagines an observer refining his initial state of belief and simply plucking out a term corresponding to the ‘data’ collected:

$$\rho = \sum_d P(d) \tilde{\rho}_d \tag{31}$$

$$\begin{array}{ccc} & \downarrow & \\ \rho & \xrightarrow{d} & \tilde{\rho}_d. \end{array} \tag{32}$$

Finally, there may be a further ‘mental readjustment’ of the observer’s beliefs, which takes into account details both of the measurement interaction and the observer’s initial quantum state. This is enacted via some (formal) unitary operation V_d :

$$\tilde{\rho}_d \rightarrow \rho_d = V_d \tilde{\rho}_d V_d^\dagger. \tag{33}$$

Putting the two processes together, one has the same result as the usual picture.

The resemblance between the process in equation (32) and the classical Bayes’ rule of equation (19) is uncanny†. By this way of viewing things, quantum collapse is indeed not such a violent state of affairs after all. Quantum measurement is nothing more, and nothing less, than a refinement and a readjustment of one’s initial state of belief. More general state changes of the form equation (14) come about similarly, but with a further step of coarse-graining (i.e. throwing away information that was in principle accessible).

Let us look at two limiting cases of efficient measurements. In the first, we imagine an observer whose initial belief structure $\rho = |\psi\rangle\langle\psi|$ is a maximally sharp state of belief. By this account, no measurement whatsoever can *refine* it. This follows because, no matter what $\{E_d\}$ is, $\rho^{1/2} E_d \rho^{1/2} = P(d) |\psi\rangle\langle\psi|$. The only state change that can come about from a measurement must be purely of the

† Other similarities between quantum collapse and Bayesian conditionalization have been discussed in [46].

mental-readjustment sort: we learn nothing new; we just change what we can predict as a consequence of the side effects of our experimental intervention. That is to say, there is a sense in which the measurement is solely disturbance. In particular, when the POVM is an orthogonal set of projectors $\{\Pi_i = |i\rangle\langle i|\}$ and the state-change mechanism is the von Neumann collapse postulate, this simply corresponds to a readjustment according to the unitary operators $U_i = |i\rangle\langle\psi|$.

At the opposite end of things, we can contemplate measurements that have no possibility at all of causing a physical disturbance to the system being measured. This could come about, for instance, by interacting with one side of an entangled pair of systems and using the consequence of that intervention to update one's beliefs about the other side. In such a case, one can show that the state change is purely of the refinement variety (with no further mental readjustment)[†]. For instance, consider a pure state $|\psi^{AB}\rangle$ whose Schmidt decomposition takes the form $|\psi^{AB}\rangle = \sum_i \lambda_i^{1/2} |a_i\rangle |b_i\rangle$. An efficient measurement on the A side of this leads to a state update of the form

$$|\psi^{AB}\rangle\langle\psi^{AB}| \longrightarrow (A_d \otimes I) |\psi^{AB}\rangle\langle\psi^{AB}| (A_d^\dagger \otimes I). \quad (34)$$

Tracing out the A side, then gives

$$\text{tr}_A \left(A_d \otimes I |\psi^{AB}\rangle\langle\psi^{AB}| A_d^\dagger \otimes I \right) = \rho^{1/2} \left(U A_d^\dagger A_d U^\dagger \right)^T \rho^{1/2}, \quad (35)$$

where ρ is the initial quantum state on the B side, U is the unitary operator connecting the $|a_i\rangle$ basis to the $|b_i\rangle$ basis, and T represents taking a transpose with respect to the $|b_i\rangle$ basis. Since the operators

$$F_d = \left(U A_d^\dagger A_d U^\dagger \right)^T \quad (36)$$

go together to form a POVM, we indeed have the claimed result.

In summary, the lesson here is that it turns out to be rather easy to think of quantum collapse as a non-commutative variant of Bayes' rule. In fact it is just in

[†] This should be contrasted with the usual picture of a 'minimally disturbing' measurement of some POVM. In our case, a minimal disturbance version of a POVM $\{E_d\}$ corresponds to taking $V_d = I$ for all d in equation (33). In the usual presentation—see [43] and [45]—it corresponds to taking $U_d = I$ for all d in equation (30) instead. For instance, Howard Wiseman writes in [45]:

The action of $[E_d^{1/2}]$ produces the minimum change in the system, required by Heisenberg's relation, to be consistent with a measurement giving the information about the state specified by the probabilities [equation (1)]. The action of $[U_d]$ represents additional back-action, an unnecessary perturbation of the system. ... A back-action evading measurement is reasonably defined by the requirement that, for all $[d]$, $[U_d]$ equals unity (up to a phase factor that can be ignored without loss of generality).

This of course means that, from the present point of view, there is no such thing as a state-independent notion of minimally disturbing measurement. Given an initial state ρ and a POVM $\{E_d\}$, the minimally disturbing measurement interaction is the one that produces pure Bayesian updating with no further (purely quantum) readjustment.

this that one starts to get a feel for a further reason for Gleason's non-contextuality assumption. In the setting of classical Bayesian conditionalization we have just that: the probability of the transition $P(h) \rightarrow P(h|d)$ is governed solely by the local probability $P(d)$. The transition does not care about how we have partitioned the rest of the potential transitions. That is, it does not care whether d is embedded in a two outcome set $\{d, \neg d\}$ or whether it is embedded in a three outcome set, $\{d, e, \neg(d \vee e)\}$, etc. Similarly with the quantum case. The probability for a transition from ρ to ρ_0 cares not whether our refinement is of the form $\rho = P(0)\rho_0 + \sum_{d=1}^{17} P(d)\rho_d$ or of the form $\rho = P(0)\rho_0 + P(18)\rho_{18}$, as long as $P(18)\rho_{18} = \sum_{d=1}^{17} P(d)\rho_d$. What could be a simpler generalization of Bayes' rule?

Indeed, leaning on that, we can restate the discussion of the 'measurement problem' at the beginning of section 4 in more technical terms. Go back to the classical setting of equation (18) where an agent has a probability distribution $P(h, d)$ over two sets of hypotheses. Marginalizing over the possibilities for d , one obtains the agent's initial belief $P(h)$ about the hypothesis h . If he gathers an explicit piece of data d , he should use Bayes' rule to update his probability about h to $P(h|d)$.

The question is this: is the transition $P(h) \rightarrow P(h|d)$ a mystery we should contend with? If someone asked for a *physical description* of this transition, would we be able to give an explanation? After all, one value for h is true and always remains true: there is no transition in it. One value for d is true and always remains true: there is no transition in it. The only discontinuous transition is in the *belief* $P(h)$, and that presumably is a property of the believer's brain. To put the issue into terms that start to sound like the quantum measurement problem, let us ask: should we not have a detailed theory of how the brain works before we can trust in the validity of Bayes' rule†?

The answer is, 'of course not!' Bayes' rule—and beyond it all of probability theory—is a tool that stands above the details of physics. Boole called probability theory a *law of thought* [48]. Its calculus specifies the optimal way an agent should reason and make decisions when faced with incomplete information. In this way, probability theory is a generalization of Aristotelian logic [41]—a tool of thought few would accept as being anchored to the details of the physical world. As far as Bayesian probability theory is concerned, a 'classical measurement' is simply any *I-know-not-what* that induces an application of Bayes' rule. It is not the task of probability theory (nor is it solvable within probability theory) to explain how the transition Bayes' rule comes about within the mind of the agent.

The formal similarities between Bayes' rule and quantum collapse may be telling us how to finally cut the Gordian knot of the measurement problem. Namely, it may be telling us that it is simply not a problem at all! Indeed, drawing on the analogies between the two theories, one is left with a spark of insight: perhaps the better part of quantum mechanics is simply 'law of thought' [49]. Perhaps the structure of the theory denotes the optimal way to reason and make decisions in light of *some* fundamental situation—a fundamental situation waiting to be ferreted out in a more satisfactory fashion.

This much we know: that fundamental situation—whatever it is—must be an ingredient Bayesian probability theory does not have. As already emphasized, there must be something to drive a wedge between the two theories. Probability

† This point was recently stated much more eloquently by Duvenhage in his paper [47].

theory alone is too general a structure. Narrowing the structure will require input from the world around us.

6.1. *Accepting quantum mechanics*

Looking at the issue from this perspective, let us ask: what does it mean to accept quantum mechanics? Does it mean accepting (in essence) the existence of an ‘expert’ whose probabilities we should strive to possess whenever we strive to be maximally rational? The key to answering this question comes from combining the previous discussion of Bayes’ rule with the considerations of the standard quantum-measurement device of section 4.2. For contemplating this will allow us to go even further than calling quantum collapse a non-commutative *variant* of Bayes’ rule.

Consider the description of quantum collapse in equations (31) through (33) in terms of one’s subjective judgements for the outcomes of a standard quantum measurement $\{E_h\}$. Using the notation there, one starts with an initial judgement $P(h) = \text{tr}(\rho E_h)$ and, after a measurement of some other observable $\{E_d\}$, ends up with a final judgement

$$P_d(h) = \text{tr}(\rho_d E_h) = \text{tr}(\tilde{\rho}_d V_d^\dagger E_h V_d) = \text{tr}(\tilde{\rho}_d F_h^d), \quad (37)$$

where $F_h^d = V_d^\dagger E_h V_d$. Note that, in general, $\{E_h\}$ and $\{E_d\}$ refer to two entirely different POVMs; the range of their indices h and d need not even be the same. Also, since $\{E_h\}$ is a minimal informationally complete POVM, $\{F_h^d\}$ will itself be informationally complete for each value of d .

Thus, *modulo a final unitary readjustment or redefinition of the standard quantum measurement* based on the data gathered, one has precisely Bayes’ rule in this transition. This follows since $\rho = \sum_d P(d) \tilde{\rho}_d$ implies

$$P(h) = \sum_d P(d) P(h|d), \quad (38)$$

with $P(h|d) = \text{tr}(\tilde{\rho}_d E_h)$.




Figure 2. A quantum measurement is any ‘I-know-not-what’ that generates an application of Bayes’ rule to one’s beliefs for the outcomes of a standard quantum measurement, that is, a decomposition of the initial state into a convex combination of other states and then a final ‘choice’ (decided by the world, not the observer) within that set. Taking into account the idea that quantum measurements are ‘invasive’ or ‘disturbing’ alters the classical Bayesian picture only in introducing a further outcome-dependent readjustment.

Another way of looking at this transition is from the ‘active’ point of view, i.e. that the axes of the probability simplex are held fixed, while the state is transformed from $P(h|d)$ to $P_d(h)$. That is, writing

$$F_h^d = \sum_{h'=1}^{D^2} \Gamma_{hh'}^d E_{h'}, \tag{39}$$

where $\Gamma_{hh'}^d$ are some real-valued coefficients and $\{E_{h'}\}$ refers to a relabelling of the original standard quantum measurement, we get

$$P_d(h) = \sum_{h'=1}^{D^2} \Gamma_{hh'}^d P(h'|d). \tag{40}$$

This gives an enticingly simple description of what quantum measurement *is* in Bayesian terms. Modulo the final readjustment, a quantum measurement is *any application of Bayes’ rule whatsoever* on the initial state $P(h)$. By any application of Bayes’ rule, I mean in particular any convex decomposition of $P(h)$ into some refinements $P(h|d)$ that also live in $\mathcal{P}_{\text{SQM}}^\dagger$. Aside from the final readjustment, a quantum measurement is just like a classical measurement: it is *any I-know-not-what* that pushes an agent to an application of Bayes’ rule \ddagger .

Accepting the formal structure of quantum mechanics is—in large part—simply accepting that it would not be in one’s best interest to hold a $P(h)$ that falls outside the convex set \mathcal{P}_{SQM} . Moreover, up to the final conditionalization rule signified by a unitary operator V_d , a measurement is simply anything that can cause an application of Bayes’ rule within \mathcal{P}_{SQM} .

But if there is nothing more than arbitrary applications of Bayes’ rule to ground the concept of quantum measurement, would not the solidity of quantum theory melt away? What else can determine when ‘this’ rather than ‘that’ measurement is performed? Surely that much has to be objective about the theory?

7. What else is information?

Suppose one wants to hold adamantly to the idea that the quantum state is purely subjective. That is, that there is no right and true quantum state for a system—the quantum state is ‘numerically additional’ to the quantum system.

\dagger Note a distinction between this way of posing Bayes’ rule and the usual way. In stating it, I give no status to a joint probability distribution $P(h, d)$. If one insists on calling the product $P(d)P(h|d)$ a joint distribution $P(h, d)$, one can do so of course, but it is only a mathematical artifice without intrinsic meaning. In particular, one should not get a feeling from $P(h, d)$ ’s mathematical existence that the random variables h and d simultaneously coexist. As always, h and d stand only for the consequences of experimental interventions into nature; without the intervention, there is no h and no d .

\ddagger Of course, I fear the wrath this phrase will bring upon me. For it will be claimed that I do not understand the first thing about the ‘problem’ of quantum measurement: it is to supply a mechanism for understanding how collapse comes about, not to dismiss it. But my language is meant to leave nothing hidden. The point here, as already emphasized in the classical case, is that it is not the task—and cannot be the task—of a theory that makes intrinsic use of probability to justify how an agent has gotten hold of a piece of information that causes him to change his beliefs. A belief is a property of one’s head, not of the object of one’s interest.

It walks through the door when the agent who is interested in the system walks through the door. Can one uphold this view, at the same time supposing *which* POVM $\{E_d\}$ and *which* state-change rule $\rho \rightarrow \rho_d = A_d \rho A_d^\dagger$ a measurement device performs are objective features of the device? The answer is no, and it is not difficult to see why.

Take as an example, a device that performs a standard von Neumann measurement $\{\Pi_d\}$, the measurement of which is accompanied by the standard collapse postulate. When a click d is found, the posterior quantum state will be $\rho_d = \Pi_d$ regardless of the initial state ρ . If this state-change rule is an objective feature of the device or its interaction with the system—i.e. it has nothing to do with the observer's subjective judgement—then the final state ρ_d too must be an objective feature of the quantum system. The argument is that simple. Moreover, it generalizes to all state change rules for which the A_d are rank-one operators without adding any further complications.

More generally, since the operators E_d control the maximal support of the final state ρ_d through $A_d = U_d E_d^{1/2}$, it must be that even the E_d are subjective judgements. For otherwise, one could say, 'Only states with support within a subspace \mathcal{S}_d are *correct*. All other states are simply *wrong*†.'

Thinking now of uninterrupted quantum time evolution as the special case of what happens to a state after the single-element POVM $\{I\}$ is performed, one is forced to the same conclusion even in that case. The time evolution super-operator for a quantum system—most generally a completely positive trace-preserving linear map on the space of operators for \mathcal{H}_D [32]—is a subjective judgement on exactly the same par as the subjectivity of the quantum state.

Here is another way of seeing the same thing. Recall what I viewed to be the most powerful argument for the quantum state's subjectivity—the Einsteinian argument of section 3. Since we can toggle the quantum state from a distance, it must not be something sitting over there, but rather something sitting over here: it can only be information about the far-away system. Let us now apply a variation of this argument to time evolutions.

Consider a simple quantum circuit on a bipartite quantum system that performs a controlled unitary operation U_i on the target. (For simplicity, let us say the bipartite system consists of two qubits.) Which unitary operation the circuit applies depends upon which state $|i\rangle$, $i = 0, 1$, of two orthogonal states impinges upon the control. Thus, for an arbitrary state $|\psi\rangle$ on the target, one finds $|i\rangle|\psi\rangle \rightarrow |i\rangle(U_i|\psi\rangle)$ for the overall evolution. Consequently the evolution of the target system alone is given by $|\psi\rangle \rightarrow U_i|\psi\rangle$. On the other hand, suppose the control is prepared in a superposition state $|\phi\rangle = \alpha|0\rangle + \beta|1\rangle$. Then the evolution for the target bit will be given by a completely positive map Φ_ϕ . That is, $|\psi\rangle \rightarrow \Phi_\phi(|\psi\rangle\langle\psi|) = |\alpha|^2 U_0|\psi\rangle\langle\psi|U_0^\dagger + |\beta|^2 U_1|\psi\rangle\langle\psi|U_1^\dagger$.

Now, to the point. Suppose rather than feeding a single qubit into the control, we feed half of an entangled pair, where the other qubit is physically far away. If an observer with this description of the set-up makes a measurement on the far-away qubit, then he will be able to induce any number of completely positive maps Φ_ϕ on the control bit. These will depend upon which measurement he performs and which outcome he gets. The point is the same as before: invoking locality, one obtains that the time evolution mapping on the single qubit cannot be an

† Such a statement is not so dissimilar to the one found in [50]. For rebuttals, see [51, 52].

objective state of affairs localized at that qubit. The time evolution, like the state, is subjective information†.

It has long been known that the trace preserving completely positive linear maps Φ over a D -dimensional vector space can be placed in a one-to-one correspondence with density operators on a D^2 -dimensional space via the relation [40]

$$\Upsilon = I \otimes \Phi(|\psi_{ME}\rangle\langle\psi_{ME}|), \tag{41}$$

where $|\psi_{ME}\rangle$ signifies a maximally entangled state on $\mathcal{H}_D \otimes \mathcal{H}_D$. This is usually treated as a representation theorem only, but maybe it is no mathematical accident. Perhaps there is a deep physical reason for it: the time evolution one ascribes to a quantum system IS a density operator! It is a state of belief no more and no less than the quantum state one assigns to the same system‡.

How to think about this? Let us go back to the issue that closed the last section. How can one possibly identify the meaning of a measurement in the Bayesian view, where a measurement ascription is itself subjective—i.e. a measurement finds a mathematical expression only in the subjective refinement of some agent’s beliefs? Here is the difficulty. When one agent contemplates viewing a piece of data d , he might be willing to use the data to refine his beliefs according to $P(h) = \sum_d P(d)P(h|d)$. However there is nothing to stop another agent from thinking the same data warrants him to refine his beliefs according to $Q(h) = \sum_d Q(d)Q(h|d)$. *A priori*, there need be no relation between the P ’s and the Q ’s.

A relation can only come from a criterion for when two agents will say that they believe they are drawing the same meaning from the data they obtain. That identification is a purely voluntary act; for there is no way for the agent to walk outside of his beliefs and see the world as it completely and totally is. The standard Bayesian solution to the problem is this: when both agents accept the same ‘statistical model’ for their expectations of the data d given a hypothesis h , then they will agree to the identity of the measurements they are each (separately) considering. That is, two agents will deem they perform the same measurement when and only when

$$P(d|h) = Q(d|h), \quad \forall h \text{ and } \forall d. \tag{42}$$

Putting this in a more evocative form, we can say that both agents agree to the meaning of a measurement when they adopt the same resolution of the identity

$$I = \sum_d \frac{P(d)P(h|d)}{P(h)} = \sum_d \frac{Q(d)Q(h|d)}{Q(h)}. \tag{43}$$

† Of course, there are sideways moves one can use to try to get around this conclusion. For instance, one could argue that, “the time evolution operator Φ on the control qubit is only an ‘effective’ evolution for it. The ‘true’ evolution for the system is the unitary evolution specified by the complete quantum circuit” [53]. In my opinion, however, moves like this are just prostrations to the Everettic temple. One could dismiss the original Einsteinian argument in the same way: “the observer toggles nothing with his localized measurement; the ‘true’ quantum state is the universal quantum state. All that is going on in a quantum measurement is the revelation of a relative state—i.e. the ‘effective’ quantum state.”

‡ Adopting this point of view sheds significant light on the old ‘problem’ of correlations that violate Bell inequalities. Explaining the issue, however, will require a separate publication.

With this, the relation to quantum measurement should be apparent. If we take it seriously that a measurement is anything that generates a refinement of one's beliefs, then an agent specifies a measurement when he specifies a resolution of his initial density operator $\rho = \sum_d P(d)\tilde{\rho}_d$. But again, there is nothing to stop another agent from thinking the data warrants a refinement that is completely unrelated to the first: $\sigma = \sum_d Q(d)\tilde{\sigma}_d$. And that is where the issue ends if the agents have no further agreement.

Just as in the classical case, however, there is a solution for the identification problem. Using the canonical construction of equation (25), we can say that both agents agree to the meaning of a measurement when they adopt the same resolution of the identity,

$$I = \sum_d P(d)\rho^{-1/2}\tilde{\rho}_d\rho^{-1/2} = \sum_d Q(d)\sigma^{-1/2}\tilde{\sigma}_d\sigma^{-1/2}. \quad (44)$$

Saying it in a more tautological way, two agents will be in agreement on the identity of a measurement when they assign it the same POVM $\{E_d\}$,

$$E_d = P(d)\rho^{-1/2}\tilde{\rho}_d\rho^{-1/2} = Q(d)\sigma^{-1/2}\tilde{\sigma}_d\sigma^{-1/2}. \quad (45)$$

The importance of this move, however, is that it draws out the proper way to think about the operators E_d from the present perspective. They play part of the role of the 'statistical model' $P(d|h)$. More generally, that role is fulfilled by the state change rule. That is to say,

$$P(d|h) \longleftrightarrow \Phi_d(\cdot) = U_d E_d^{1/2} \cdot E_d^{1/2} U_d^\dagger. \quad (46)$$

The completely positive map that gives a mathematical description to quantum time evolution is just such a map. Its role is that of the subjective statistical model $P(d|h)$, where d happens to be drawn from a one-element set. Thus, thinking back on entanglement, it seems the general structure of quantum time evolutions cannot be the wedge we are looking for either. What we see instead is that there is a secret waiting to be unlocked, and when it is unlocked, it will likely tell us as much about quantum time evolutions as quantum states and quantum measurements.

8. Intermission

Until now I have tried to tear down as much of quantum mechanics as possible. Section 3 argued that quantum states—whatever they be—cannot be objective entities. Section 4 argued that there is nothing sacred about the quantum probability rule and the best way to think of a quantum state is as a state of belief about what *would* happen if one *were* to approach a standard measurement device. Section 5 argued that even quantum entanglement is a secondary and subjective effect. Section 6 argued that all a measurement *is* is just an arbitrary application of Bayes' rule—an arbitrary refinement of one's beliefs—along with some account that measurements are invasive interventions into nature. Section 7 argued that even quantum time evolutions are subjective judgements; they just so happen to be conditional judgements. . . . And so it went.

Subjective. Subjective! Subjective!! It is a word that will not go away. But finding overwhelming subjectivity in the theory is not something to be

proud of. There are limits: the last thing we need is a bloodbath of deconstruction. At the end of the day, there had better be some term, some element in quantum theory that stands for the objective, or we might as well melt away and call the whole world a dream.

9. The oyster and the quantum

A grain of sand falls into the shell of an oyster and the result is a pearl. The oyster's sensitivity to the touch is the source of a beautiful gem. In the 75 years that have passed since the founding of quantum mechanics, only the last 10 have turned to a view and an attitude that may finally reveal the essence of the theory. The quantum world is sensitive to the touch, and that may be one of the best things about it. Quantum information—with its three specializations of quantum information theory, quantum cryptography and quantum computing—leads the way in telling us how to quantify this idea. Quantum algorithms can be exponentially faster than classical algorithms. Secret keys can be encoded into physical systems in such a way as to reveal whether information has been gathered about them. The list of triumphs keeps growing.

The key to so much of this has been simply in a change of attitude. This can be seen by going to almost any older textbook on quantum mechanics: nine times out of ten, the Heisenberg uncertainty relation is presented in a way that conveys the feeling that we have been short-changed by the physical world. 'Look at classical physics, how nice it is: we can measure a particle's position and momentum with as much accuracy as we wish. How limiting quantum theory is instead. We have $\Delta x \Delta p \geq \hbar/2$, and there is nothing we can do about it. The task of physics is to just sober up to this and make the best of it.'

How this contrasts with the point of departure of quantum information! There the task is not to ask what limits quantum mechanics places upon us, but what novel, productive things we can do in the quantum world that we could not have done otherwise. In what ways is the quantum world fantastically better than the classical one?

If one is looking for something 'real' in quantum theory, what more direct tack could one take than to look to its technologies? People may argue about the objective reality of the wave function *ad infinitum*, but few would argue about the existence of quantum cryptography as a solid prediction of the theory. Why not take that or a similar effect as the grounding for what quantum mechanics is trying to tell us about nature?

Let us give this imprecise set of thoughts some shape by re-expressing quantum cryptography in the language built up in the previous sections. For quantum key distribution it is essential to be able to prepare a physical system in one or another quantum state drawn from some fixed *non-orthogonal* set [54]. These non-orthogonal states are used to encode a potentially secret cryptographic key to be shared between the sender and receiver. The information an eavesdropper seeks is about which quantum state was actually prepared in each individual transmission. What is novel here is that the encoding of the proposed key into non-orthogonal states forces the information-gathering process to induce a disturbance to the overall *set* of states. That is, the presence of an active eavesdropper transforms the initial pure states into a set of mixed states or, at the very least, into a set of pure states with larger overlaps than before. This action ultimately boils down to a loss

of predictability for the sender over the outcomes of the receiver's measurements and, so, is directly detectable by the receiver (who reveals some of those outcomes for the sender's inspection). More importantly, there is a direct connection between the statistical information gained by an eavesdropper and the consequent disturbance she must induce to the quantum states in the process. As the information gathered goes up, the necessary disturbance also goes up in a precise way [55].

Note the two ingredients that appear in this. First, the information gathering or measurement is grounded with respect to one observer (in this case, the eavesdropper), while the disturbance is grounded with respect to another (here, the sender). In particular, the disturbance is to the sender's previous information—this is measured by her diminished ability to predict the outcomes of certain measurements the legitimate receiver might perform. No hint of any variable intrinsic to the system is made use of in this formulation of the idea of 'measurement causing disturbance'.

The second ingredient is that one must consider at least two non-orthogonal preparations for the formulation to have any meaning. This is because the information gathering is not about some classically-defined observable—i.e. about some unknown hidden variable or reality intrinsic to the system—but instead about which unknown state the sender actually prepared. The lesson is this: forget about the unknown preparation, and the random outcome of the measurement is information about nothing. It is simply 'quantum noise' with no connection to any pre-existing variable.

How crucial is this second ingredient—that is, that there be at least two non-orthogonal states within the set under consideration? We can address its necessity by making a shift in the account above: one might say that the eavesdropper's goal is not so much to uncover the identity of the unknown quantum state, but to sharpen her predictability over the receiver's measurement outcomes. In fact, she would like to do this at the same time as disturbing the sender's predictions as little as possible. Changing the language still further to the terminology of section 4, the eavesdropper's actions serve to sharpen her information about the potential consequences of the receiver's further interventions on the system. (Again, she would like to do this while minimally diminishing the sender's previous information about those same consequences.) In the cryptographic context, a by-product of this effort is that the eavesdropper ultimately comes to a more sound prediction of the secret key. From the present point of view, however, the importance of this change of language is that it leads to an almost Bayesian perspective on the information-disturbance problem.

Within Bayesian probability the most significant theme is to identify conditions under which a set of decision-making agents can come to a common probability assignment for some random variable despite their initial differences [41]. One might similarly view the process of quantum eavesdropping. The sender and the eavesdropper start off initially with differing quantum state assignments for a single physical system. In this case it so happens that the sender can make sharper predictions than the eavesdropper about the outcomes of the receiver's measurements. The eavesdropper, not satisfied with the situation, performs a measurement on the system to sharpen those predictions. In particular, there is an attempt to come into something of an agreement with the sender but without revealing the outcomes of her measurements or her very presence.

At this point a distinct *property* of the quantum world makes itself known. The eavesdropper's attempt to surreptitiously come into alignment with the sender's predictability is always shunted away from its goal. This shunting of various observer's predictability is the subtle manner in which the quantum world is sensitive to our experimental interventions.

Maybe this is our crucial hint! The wedge that drives a distinction between Bayesian probability theory in general and quantum mechanics in particular is perhaps nothing more than this 'Zing!' of a quantum system that is manifested when an agent interacts with it. It is this wild sensitivity to the touch that keeps our information and beliefs from ever coming into too great an alignment. The most our beliefs about the consequences of our interventions on a system can come into alignment is captured by the mathematical structure of a pure quantum state $|\psi\rangle$. Take all possible information–disturbance curves for a quantum system, tie them into a bundle, and *that* is the long-awaited property, the input we have been looking for from nature. Or, that is the speculation.

How might one hope to mathematize the bundle of all possible information–disturbance curves for a system? If it can be done at all, the effort will have to end up depending upon a single real parameter—the dimension of the system's Hilbert space†. As a safety check, let us ask ourselves whether this is a tenable possibility? Or will the Hilbert-space dimension go the wayside of subjectivity, just as so many of the other terms in the theory? I think the answer will be in the negative: the Hilbert-space dimension will survive to be a stand-alone concept with no need of an agent for its definition.

The simplest check perhaps is to pose the Einsteinian test for it as we did first for the quantum state and then for quantum time evolutions. Posit a bipartite system with Hilbert spaces \mathcal{H}_{D_1} and \mathcal{H}_{D_2} (with dimensions D_1 and D_2 respectively) and imagine an initial quantum state for that bipartite system. As argued too many times already, the quantum state must be a subjective component in the theory because the theory allows localized measurements on the D_1 system to change the quantum state for the D_2 system. In contrast, is there anything one can do at the D_1 site to change the numerical value of D_2 ? It does not appear so. Indeed, the only way to change that number is to scrap the initial supposition. Thus, to this extent, one has every right to call the numbers D_1 and D_2 potential 'elements of reality'.

It may not look like much, but it is a start. And one should not underestimate the power of a hint, no matter how small.

Acknowledgments

I thank C. Caves, G. Comer, D. Mermin and R. Schack for years of correspondence, J. Bub and L. Hardy for courage, and J. Nicholson for the paper's figures. I also thank M. Appleby, H. Barnum, C. Bennett, H. Bernstein, A. Cabello, M. Donald, S. van Enk, J. Finkelstein, P. Grangier, O. Hirota, A. Landahl, H. Mabuchi, J. Malley, U. Mohrhoff, M. Nielsen, M. Ozawa, A. Peres, A. Plotnitsky, J. Preskill, T. Rudolph, J. Summhammer, C. Timpson and A. Wilce for comments.

† The importance of this single parameter has also been pointed out in [56].

References

- [1] FUCHS, C. A., 2002, quant-ph/0205039.
- [2] CUSHING, J. T., FINE, A., and GOLDSTEIN, S. (eds), 1996, *Bohmian Mechanics and Quantum Theory: An Appraisal* (Dordrecht: Kluwer).
- [3] GRIFFITHS, R. B., and OMNÈS, R., 1999, *Phys. Today*, **52**, 26.
- [4] CRAMER, J. G., 1988, *Int. J. theor. Phys.*, **27**, 227.
- [5] GHIRARDI, G. C., and PEARLE, P., 1990, *PSA 1990*, edited by A. Fine, M. Forbes and L. Wessels (East Lansing: Philosophy of Science Association), p. 19.
- [6] ZUREK, W. H., 1996, *Phil. Trans. R. Soc. (London) A*, **356**, 1793.
- [7] GRANGIER, P., 2000, quant-ph/0012122; 2001, quant-ph/0111154.
- [8] DEUTSCH, D., 1997, *The Fabric of Reality* (New York: Allen Lane).
- [9] ROVELLI, C., 1996, *Int. J. theor. Phys.*, **35**, 1637.
- [10] MISNER, C. W., THORNE, K. S., and WHEELER, J. A., 1973, *Gravitation* (San Francisco: Freeman), p. 23.
- [11] ZEILINGER, A., 1999, *Found. Phys.*, **29**, 631; BRUKNER, C., and ZEILINGER, A., 1999, *Phys. Rev. Lett.*, **83**, 3354; DEVETAK, I., and STAPLES, A. E., 2001, quant-ph/0112166; BARNUM, H., 2002, quant-ph/0205129.
- [12] GRANGIER, P., 2001, personal communication.
- [13] BENNETT, C. H., BRUN, T. A., CAVES, C. M., and MERMIN, N. D., 2001, Various vibes.
- [14] BENNETT, C. H., BRASSARD, G., CRÉPEAU, C., JOZSA, R., PERES, A., and WOOTTERS, W. K., 1993, *Phys. Rev. Lett.*, **70**, 1895.
- [15] BELL, J. S., 1987, *Speakable and Unsayable in Quantum Mechanics* (Cambridge: Cambridge University Press).
- [16] NIELSEN, M. A., and CHUANG, I. L., 2000, *Quantum Computation and Quantum Information*, (Cambridge: Cambridge University Press).
- [17] BUSCH, P., 1999, quant-ph/9909073.
- [18] CAVES, C. M., FUCHS, C. A., MANNE, K., and RENES, J., 2000, unpublished.
- [19] GLEASON, A. M., 1957, *J. math. Mech.* **6**, 885.
- [20] EINSTEIN, A., PODOLSKY, B., and ROSEN, N., 1935, *Phys. Rev.*, **47**, 777.
- [21] FINE, A., 1986, *The Shaky Game: Einstein, Realism and the Quantum Theory* (Chicago: University Chicago Press).
- [22] BERNSTEIN, J., 1991, *Quantum Profiles* (Princeton: Princeton University Press).
- [23] EINSTEIN, A., 1970, *Albert Einstein: Philosopher-Scientist*, edited by P. A. Schilpp (La Salle: Open Court), p. 665.
- [24] PERES, A., 1993, *Quantum Theory: Concepts and Methods* (Dordrecht: Kluwer).
- [25] CAVES, C. M., and FUCHS, C. A., 1996, *Ann. Israel phys. Soc.*, **12**, 226.
- [26] PAULI, W., 1988, *Beyond the Atom: The Philosophical Thought of Wolfgang Pauli*, edited by K. V. Laurikainen (Berlin: Springer-Verlag), p. 226.
- [27] PENROSE, R., 1989, *The Emperor's New Mind* (Oxford: Oxford University Press), p. 268.
- [28] SAVAGE, L. J., 1954, *The Foundations of Statistics* (New York: Dover).
- [29] FUCHS, C. A., and PERES, A., 2000, *Phys. Today*, **53**, 70; FUCHS, C. A., and PERES, A., 2000, *Phys. Today*, **53**, 14.
- [30] BARNUM, H., CAVES, C. M., FINKELSTEIN, J., FUCHS, C. A., and SCHACK, R., 2000, *Proc. Roy. Soc. (London) A*, **456**, 1175.
- [31] COOKE, R., KEANE, M., and MORAN, W., 1981, *Math. Proc. Camb. phil. Soc.*, **98**, 117; PITOWSKY, I., 1998, *J. math. Phys.*, **39**, 218.
- [32] KRAUS, K., 1983, *States, Effects, and Operations* (Berlin: Springer).
- [33] MEYER, D. A., 1999, *Phys. Rev. Lett.*, **83**, 3751; KENT, A., 1999, *Phys. Rev. Lett.*, **83**, 3755.
- [34] HARDY, L., 2001, quant-ph/0101012; 2001, quant-ph/0111068.
- [35] PRUGOVEČKI, E., 1977, *Int. J. theor. Phys.*, **16**, 321; SCHROECK JR, F. E., 1989, *Int. J. theor. Phys.*, **28**, 247; BUSCH, P., 1991, *Int. J. theor. Phys.*, **30**, 1217.
- [36] CAVES, C. M., FUCHS, C. A., and SCHACK, R., 2001, quant-ph/0104088.
- [37] PERES, A., and TERNO, D., 1998, *J. Phys. A*, **31**, L671.
- [38] WALLACH, N. R., 2000, quant-ph/0002058.

- [39] ARAKI, H., 1980, *Commun. math. Phys.*, **75**, 1; WOOTTERS, W. K., 1986, *Found. Phys.*, **16**, 391.
- [40] JAMIOLKOWSKI, A., 1972, *Rep. math. Phys.*, **3**, 275; CHOI, M.-D., 1975, *Lin. Alg. Appl.*, **10**, 285.
- [41] KYBURG JR, H. E., and SMOKLER, H. E. (eds), 1980, *Studies in Subjective Probability* (Huntington: Krieger); BERNARDO, J. M., and SMITH, A. F. M., 1994, *Bayesian Theory* (Chichester: Wiley).
- [42] WEHRL, A., 1978, *Rev. mod. Phys.*, **50**, 221.
- [43] FUCHS, C. A., and JACOBS, K., 2001, *Phys. Rev. A*, **63**, 062305.
- [44] NIELSEN, M. A., 2001, *Phys. Rev. A*, **63**, 022114.
- [45] WISEMAN, H. M., 1995, *Phys. Rev. A*, **51**, 2459.
- [46] OZAWA, M., 1997, *Ann. Phys.*, **259**, 121; OZAWA, M., 1998, *Fortschr. Phys.*, **46**, 615.
- [47] DUVENHAGE, R., 2002, quant-ph/0203070.
- [48] BOOLE, G., 1958, *An Investigation of the Laws of Thought* (New York: Dover).
- [49] FUCHS, C. A., 2001, quant-ph/0105039.
- [50] BRUN, T. A., FINKELSTEIN, J., and MERMIN, N. D., 2002, *Phys. Rev. A*, **65**, 032315.
- [51] FUCHS, C. A., *Quantum States: W.H.A.T.?*, posted <http://netlib.bell-labs.com/who/cafuchs/>.
- [52] CAVES, C. M., FUCHS, C. A., and SCHACK, R., 2002, quant-ph/0206110.
- [53] CAVES, C. M., 1999, personal communication.
- [54] BENNETT, C. H., BRASSARD, G., and MERMIN, N. D., 1992, *Phys. Rev. Lett.*, **68**, 557; FUCHS, C. A., 1998, *Fort. Phys.*, **46**, 535.
- [55] FUCHS, C. A., and PERES, A., 1996, *Phys. Rev. A*, **53**, 2038; FUCHS, C. A., GISIN, N., GRIFFITHS, R. B., NIU, C.-S., and PERES, A., 1997, *Phys. Rev. A*, **56**, 163; BRUSS, D., 1998, *Phys. Rev. Lett.*, **81**, 2598; BANASZEK, K., and DEVETAK, I., 2001, *Phys. Rev. A*, **64**, 052307.
- [56] BLUME-KOHOUT, R., CAVES, C. M., and DEUTSCH, I. H., 2002, quant-ph/0204157.