

# “It from bit” and the quantum probability rule

M. S. Leifer

June 16, 2013

## Abstract

I argue that, on the subjective Bayesian interpretation of probability, “it from bit” requires a generalization of probability theory. This does not get us all the way to the quantum probability rule because an extra constraint, known as noncontextuality, is required. I outline the prospects for a derivation of noncontextuality within this approach and argue that it requires a realist approach to physics, or “bit from it”. I then explain why this does not conflict with “it from bit”.

## 1 Wheeler’s “it from bit”

It from bit. Otherwise put, every it—every particle, every field of force, even the spacetime continuum itself—derives its function, its meaning, its very existence entirely—even if in some contexts indirectly—from the apparatus-elicited answers to yes or no questions, binary choices, bits.

It from bit symbolizes the idea that every item of the physical world has at bottom—at a very deep bottom, in most instances—an immaterial source and explanation; that what we call reality arises in the last analysis from the posing of yes-no questions and the registering of equipment-evoked responses; in short, that all things physical are information-theoretic in origin and this is a participatory universe.

— J. A. Wheeler [1]

John Wheeler’s “it from bit” is a thesis about the foundations of quantum theory. It says that the things that we usually think of as real—particles, fields and even spacetime—have no existence independent of the questions that we ask about them. When a detector clicks it is not registering something that was there independently of the experiment. Rather, the very act of setting up the detector in a certain way—the choice of question—is responsible for the occurrence of the click. It is only the act of asking questions that causes the answers to come into being. This idea is perhaps best illustrated by Wheeler’s parable of the game of twenty questions (surprise version).

You recall how it goes—one of the after-dinner party sent out of the living room, the others agreeing on a word, the one fated to be a questioner returning and starting his questions. “Is it a living object?” “No.” “Is it here on earth?” “Yes.” So the questions go from respondent to respondent around the room until at length the word emerges: victory if in twenty tries or less; otherwise, defeat.

Then comes the moment when we are fourth to be sent from the room. We are locked out unbelievably long. On finally being readmitted, we find a smile on everyone’s face, sign of a joke or a plot. We innocently start our questions. At first the answers come quickly. Then each question begins to take longer in the answering—strange, when the answer itself is only a simple “yes” or “no.” At length, feeling hot on the trail, we ask, “Is the word ‘cloud’?” “Yes,” comes the reply, and everyone bursts out laughing. When we were out of the room, they explain, they had agreed not to agree in advance on any word at all. Each one around the circle could respond “yes” or “no” as he pleased to whatever question we put to him. But however he replied he had

to have a word in mind compatible with his own reply—and with all the replies that went before.  
No wonder some of those decisions between “yes” and “no” proved so hard!

— J. A. Wheeler [2]

Wheeler proposed “it from bit” as a clue to help us answer the question “How come the quantum?”, i.e. to derive the mathematical apparatus of quantum theory from a set of clear physical principles. In this essay, I discuss whether “it from bit” implies the quantum probability rule, otherwise known as the Born rule, which would get us part of the way towards answering Wheeler’s question.

My main argument is that, on the subjective Bayesian interpretation of probability, “it from bit” requires a generalized probability theory. I explain why this is not ruled out by the common claim that classical probability theory is not to be violated on pain of irrationality.

In the context of quantum theory, “it from bit” does not quite get us all the way to the Born rule because the latter mandates a further constraint known as noncontextuality. The prospects for understanding noncontextuality as a rationality requirement or an empirical addition are slim. Extra physical principles are needed and I argue that these must be about the nature of reality, rather than the nature of our knowledge. This seems to conflict with “it from bit” as it requires an agent-independent reality, suggesting “bit from it”. I argue that there is no such conflict because the sense of “it” used in “it from bit” is different from the sense used in “bit from it”.

## 2 The interpretation of probability

Since von Neumann’s work on quantum logic and operator algebras [3, 4], it has been known that quantum theory can be viewed as a generalization of probability theory [5, 6]. If we want to understand what this tells us about the nature of reality then we will need to adopt a concrete theory of how probabilities relate to the world, which is the job of an interpretation of probability theory<sup>1</sup>. Three main classes of interpretation have arisen to meet this need: frequentism (probability is long run relative frequency), epistemic probability (probabilities represent the knowledge, information, or beliefs of a decision making agent), and objective chance (probabilities represent a kind of law of nature or a disposition for a system to act in a certain way). Getting into the details of all these options would take us too far afield, but a few comments are in order to explain why adopting my preferred epistemic interpretation, known as subjective Bayesianism<sup>2</sup>, is not a crazy thing to do.

Frequentism is still popular amongst physicists, but it has largely been abandoned by scholars of the philosophy of probability. It is not able to handle single-case probabilities, e.g. the probability that civilization will be destroyed by a nuclear war, and it leads to a bizarre reading of the law of large numbers that does not do the explanatory work required of it<sup>3</sup>. A common position in the philosophy of probability is that subjective Bayesianism is more satisfactory, but that it needs to be backed up by some theory of objective chance in order to account for probabilistic laws<sup>4</sup>. My own view is that subjective Bayesianism suffices on its own, but whether or not one believes in objective chance is irrelevant for the present discussion, since objective chances need to be connected to epistemic probabilities in some way in order to explain how we can come to know statistical laws. The usual way of doing this is via David Lewis’ principal principle [13]. One of the implications of this is that objective chances must have the same mathematical structure as subjective Bayesian probabilities. Therefore, if we can argue that “it from bit” requires a modification of subjective Bayesian probability then the same will apply to objective chances as well. It is also worth noting that several modern interpretations of quantum theory adopt subjective Bayesianism, including “Quantum Bayesianism” [14, 15, 16, 17] and the Deutsch-Wallace variant of many-worlds [18, 19, 20, 21] amongst others [22].

---

<sup>1</sup>See [7] for an accessible introduction and [8] for a collection of key papers.

<sup>2</sup>Subjective Bayesianism has its origins in [9, 10]. An accessible introduction is [11].

<sup>3</sup>See [12] for a critique of frequentism in statistics.

<sup>4</sup>This view originates with David Lewis [13].

### 3 Subjective Bayesian probability

Subjective Bayesianism says that probabilities represent the degrees of belief of a decision making agent, who is conventionally described in the second person as “you”. Degrees of belief are measured by looking at your behaviour, e.g. your willingness to enter into bets. The claim is that if you do not structure your beliefs according to the axioms of probability theory then you are irrational. There are various ways of deriving this, differing in their simplicity and sophistication. For ease of exposition, I base my discussion on the simplest approach, known as the Dutch book argument.

The Dutch book argument defines your degree of belief in the occurrence of an uncertain event  $E$  as the value  $\$p(E)$  you consider to be a fair price for a lottery ticket that pays \$1 if  $E$  occurs and nothing if it does not. “Fair price” here means that you would be prepared to buy or sell any number of these tickets at that price and that you would be prepared to do this in combination with fair bets on arbitrary sets of other events. Your degrees of belief are said to be irrational if a malicious bookmaker can force you to enter into a system of bets that would cause you to lose money whatever the outcome, despite the fact that you consider them all fair. Otherwise, your degrees of belief are said to be rational. The Dutch book argument then shows that your degrees of belief are rational if, and only if, they satisfy the usual axioms of probability theory. These axioms are:

- **Background framework:** There is a set  $\Omega$ , called the sample space, containing the most fine-grained events you might be asked to bet on, e.g. if you are betting on the outcome of a dice roll then  $\Omega = \{1, 2, 3, 4, 5, 6\}$ . In general, an event is a subset of  $\Omega$ , e.g. the event that the dice roll comes out odd is  $\{1, 3, 5\}$ . For simplicity, we assume that  $\Omega$  is finite. The set of events forms a Boolean algebra, which just means that it supports the usual logical notions of AND, OR and NOT.
- **A1:** For all events  $E \subseteq \Omega$ ,  $0 \leq p(E) \leq 1$ .
- **A2:** For the certain event  $\Omega$ ,  $p(\Omega) = 1$ .
- **A3:** If  $E \cap F = \emptyset$ , i.e.  $E$  and  $F$  cannot both happen together, then  $p(E \cup F) = p(E) + p(F)$ , where  $E \cup F$  means the event that either  $E$  or  $F$  occurs.

For illustrative purposes, here is the part of the argument showing that violations of **A1** and **A2** are irrational. Consider an event  $E$  and suppose contra **A1** that  $p(E) < 0$ . This means that you would be willing to sell a lottery ticket that pays out on  $E$  to the bookie for a negative amount of money, i.e. you would pay her  $\$p(E)$  to take the ticket off your hands. Now, if  $E$  occurs you will have to pay the bookie \$1 so in total you will have paid her  $\$1 + p(E)$ , and if  $E$  does not occur you will have paid her a total of  $\$p(E)$ . Either way, you will lose money so having negative degrees of belief is irrational. A similar argument shows that having degrees of belief larger than 1 is irrational. Now suppose contra **A2** that  $p(\Omega) < 1$ . Then, you would be prepared to sell the lottery ticket for  $\$p(\Omega)$  and pay out \$1 if  $\Omega$  occurs. However, since  $\Omega$  is certain to occur, you will always end up paying out, which leaves you with a loss.

### 4 Is probability theory normative?

Based on this kind of argument, many subjective Bayesians regard probability theory as akin to propositional logic<sup>5</sup>. In logic, you start with a set of premises that you regard as true, and then you use the rules of logic to figure out what other propositions must be true or false as a consequence. If you fail to abide by those truth values then there is an inconsistency in your reasoning. However, there is nothing in logic that tells you what premises you have to start with. The premises are simply the input to the problem and logic tells you what else they compel you to believe. Similarly, subjective probability does not tell you what numerical value you must assign to any uncertain event<sup>6</sup>, but given some of those values as input, it tells you what

<sup>5</sup>For example, see [23] where this argument is made repeatedly.

<sup>6</sup>This is where it differs from objective Bayesianism [24], which asserts that there is a unique rational probability that you ought to assign. However, defining such a unique probability is problematic at best.

values you must assign to other events on pain of inconsistency, the inconsistency here being exposed in the form of a sure loss. Like logic, subjective Bayesians regard probability theory as normative rather than descriptive, i.e. they claim that you should structure your degrees of belief about uncertain events according to probability theory if you aspire to be ideally rational, but not that humans actually do structure their beliefs in this way. In fact, much research shows that they do not [25].

The normative view of probability theory presents a problem if we want to view quantum theory as generalized probability because it implies that it is irrational to use anything other than conventional probability theory to reason about uncertain events. Fortunately, the normative view is not just wrong, but obviously wrong. Unlike logic, it is easy to come up with situations in which the Dutch book argument has no normative force. Because of this, the idea that it might happen in quantum theory too is not particularly radical.

For example, the Dutch book argument requires that you view the fair price for selling a lottery ticket to be the same as the fair price for buying it. In reality, people are more reluctant to take on risk than they are to maintain a risk for which they have already paid the cost. Therefore, the fair selling price might be higher than the fair buying price. This leads to the more general theory of upper and lower probabilities wherein degrees of belief are represented by intervals on the real line rather than precise numerical values [26].

At this point, I should address the fact that the Dutch book argument is not the only subjective Bayesian derivation of probability theory, so its defects may not be shared with the other derivations. The most general subjective arguments for probability theory are formulated in the context of decision theory, with Savage's axioms being the most prominent example [27]. These take account of things like the fact that you may be risk averse and your appreciation of money is nonlinear, e.g. \$1 is worth more to a homeless person than a billionaire, so they replace the financial considerations of the Dutch book argument with the more general concept of "utility". However, what all these arguments have in common is that they are hedging strategies. They start with some set of uncertain events and then they introduce various decision scenarios that you could be faced with where the consequences depend on uncertain events in some way, e.g. the prizes in a game that depends on dice rolls. Importantly, these arguments only work if the set of decision scenarios is rich enough. They ask you to consider situations in which the prizes for the various outcomes are chopped up and exchanged with each other in various ways. For example, in the Dutch book argument this comes in the form of the idea that you must be prepared to buy or sell arbitrarily many tickets for arbitrary sets of events at the fair price. The arguments then conclude that if you do not structure your beliefs according to probability theory then there is some decision scenario in which you would be better off had you done so and none in which you would be worse off. However, in real life, it is rather implausible that you would be faced with such a rich set of decision scenarios. More often, you know something in advance about what decisions you are going to be faced with. This is why decision theoretic arguments are hedging strategies. They start from a situation in which you do not know what decisions you are going to be faced with and then they ask you to consider the worst possible scenario. If you know for sure that this scenario is not going to come up then the arguments have no normative force.

As an example, consider the following scenario. There is a coin that is going to be flipped exactly once. You have in your possession \$1 and you are going to be forced to bet that dollar on whether the coin will come up heads or tails, with a prize of \$2 if you get it right. You do not have the option of not placing a bet. How should you structure your beliefs about whether the coin will come up heads or tails? If the decision theoretic arguments applied then we would be forced to say that you must come up with a precise numerical value for the probability of heads  $p(H)$ . However, it is clear that the cogitation involved in coming up with this number is completely pointless in this scenario. All you need to know is the answer to a single question. Do you think heads is more likely to come up than tails? Your decision is completely determined by this answer, which is just a single bit of information rather than a precise numerical value.

It should be clear from this that the decision theoretic arguments are not as strongly normative as the laws of logic. Instead they are *conditionally* normative, i.e. normative if the decision scenarios envisaged in the argument are all possible.

## 5 “It from bit” implies a generalized probability theory

A universe that obeys “it from bit” is a universe in which not all conceivable decision scenarios are possible. To explain this, consider again Wheeler’s parable of twenty questions (surprise version) and imagine that you are observing the game passively, placing bets with a bookmaker on the side as it proceeds. To make things more analogous to quantum theory, imagine that the respondents exit the room as soon as they have answered their question, never to be heard from again. We might imagine that they are sent through a wormhole into a region of spacetime that will forever be outside of our observable universe and that the wormhole promptly closes as soon as they enter it. This rules out the possibility that we might ask them about what they would have answered if they had been asked a different question, since in quantum theory we generally cannot find out what the outcome of a measurement that we did not actually make would have been.

Suppose that, at some point in the game, you make a bet with the bookie that the object that the fifth respondent has in mind is a dove. However, what actually happens is that the questioner asks “Is it white?” and the answer comes back “yes”, whereupon the fifth respondent is whisked off to the far corners of the universe. Now, although the answer “yes” is consistent with the object being a dove, this is not enough to resolve the bet as there are plenty of other conceivable white objects. As in Wheeler’s story, suppose that the last question asked is “Is it a cloud?” and that the answer comes back “yes”. In the usual version of twenty questions this would be enough to resolve the bet in the bookie’s favor because all the respondents are thinking of a common object. However, in the surprise version this is not the case. It could well be that “dove” was consistent with all the answers given so far at the time we made the bet, and that the fifth respondent was actually thinking of a dove. We can never know and so the bet can never be resolved. It has to be called off and you should get a refund.

Whilst the bet described above is unresolvable, other bets are still jointly resolvable, e.g. a bet on whether the fifth respondent was thinking of a white object together with a bet on whether the last respondent was thinking of a cloud. The set of bets that is jointly resolvable depends on the sequence of questions that is actually asked by the questioner. If you want to develop a hedging strategy ahead of time, then you need to consider all possible sequences of questions that might be asked to ensure that you cannot be forced into a sure loss for any of them.

For the subjective Bayesian, the main lesson of this is that, in general, only certain subsets of all possible bets are jointly resolvable. Define a *betting context* to be a set of events such that bets on all of them are jointly resolvable and to which no other event can be added without violating this condition. It is safe to assume that each betting context is a Boolean algebra, since, if we can find out whether  $E$  occurred at the same time as finding out whether  $F$  occurred, then we can also determine whether they both occurred, whether either one of them occurred, and whether they failed to occur, so we can define the usual logical notions of AND, OR and NOT. However, unlike in conventional probability theory, there need not be a common algebra on which all of the events that occur in different betting contexts are jointly defined. Because of this, the Dutch book argument has normative force within a betting context, but it does not tell us how probabilities should be related across different contexts. Therefore, our degrees of belief should be represented by a set of probability distributions  $p(E|\mathcal{B})$ , one for each betting context  $\mathcal{B}$ <sup>7</sup>.

This framework can be applied to quantum theory where the betting contexts represent sets of measurements that can be performed together at the same time. The details of this are rather technical, so they are relegated to appendix A. The probabilities that result from this are more general than those allowed by quantum theory. To get uniquely down to the Born rule, we need an extra constraint, known as *noncontextuality*. This says that there are certain pairs of events from different betting contexts,  $E \in \mathcal{B}$  and  $F \in \mathcal{B}'$ , that must always be assigned the same probability  $p(E|\mathcal{B}) = p(F|\mathcal{B}')$ . Therefore, we need to explain how such additional constraints can be understood.

---

<sup>7</sup>Despite the notation,  $p(E|\mathcal{B})$  is not a conditional probability distribution because there need not be a common algebra on which all the events are defined. Some authors do not consider this to be a generalization of probability theory [28, 21], since all we are saying is that we have a bunch of probability distributions rather than just one. However, such systems can display nonclassical features such as violations of Bell inequalities and no-cloning [29] so they are worthy of the name “generalization” if anything is.

## 6 Noncontextuality in subjective Bayesianism

One option is that noncontextuality could simply be posited as an additional fundamental principle. Previous Dutch book arguments for the Born rule have done essentially this [14, 22]. However, subjective Bayesians do not accept fundamental constraints on probabilities beyond those required by rationality. Imposing such constraints would be like saying that you are allowed to construct a logical argument providing one of your starting premises is “the car is red”, but if you start from “the car is yellow” then any argument you make is logically invalid. Additional constraints on probabilities are contingent facts about your state of belief, just as logical premises are contingent facts about the world. Therefore, noncontextuality needs to be derived in some way.

One possibility is that noncontextuality follows from logical equivalence, i.e. if quantum theory always assigns the same probability to  $E$  in context  $\mathcal{B}$  and  $F$  in  $\mathcal{B}'$  then these should be regarded as equivalent logical statements, in the same sense that  $E$  and NOT(NOTE) are equivalent in a Boolean algebra<sup>8</sup>. Logical equivalence implies that it ought to be possible to construct a Dutch book that results in a sure loss if  $p(E|\mathcal{B}) \neq p(F|\mathcal{B}')$ . This can only be done if you are willing to accept that the occurrence of  $E$  in betting context  $\mathcal{B}$  makes it necessary that  $F$  would have occurred had the betting context been  $\mathcal{B}'$  and vice versa. If this is the case, then you will agree that a bet made on  $F$  in betting context  $\mathcal{B}'$  should also pay out if the betting context was in fact  $\mathcal{B}$  and the event  $E$  occurred and vice versa. If this is the case, then the bookie can construct a Dutch book against  $p(E|\mathcal{B}) < p(F|\mathcal{B}')$  by buying a ticket from you that pays out on  $E$  and selling a ticket that pays out on  $F$ . The payouts on these tickets will be the same, so you will lose money in this transaction. By exchanging the roles of  $E$  and  $F$ , there would be a Dutch book against  $p(E|\mathcal{B}) > p(F|\mathcal{B}')$  as well.

This strategy hinges on whether it is reasonable to make counterfactual assertions, i.e. assertions about what would have happened had the betting context been different. However, “it from bit” declares such counterfactuals meaningless because it says that there is no answer to questions that have not been asked. Even if we do not accept “it from bit”, the Kochen-Specker theorem [30] implies that counterfactual assertions cannot all respect noncontextuality, i.e. there would have to be pairs of events  $E \in \mathcal{B}$  and  $F \in \mathcal{B}'$  such that if  $E$  occurs in  $\mathcal{B}$  then  $F$  would not have occurred in  $\mathcal{B}'$  even though quantum theory asserts that  $p(E|\mathcal{B}) = p(F|\mathcal{B}')$  always holds. We conclude that noncontextuality of probability assignments cannot be a rationality requirement.

Another possibility is that noncontextuality could be adopted simply because we have performed many quantum experiments and have always observed relative frequencies in accord with the Born rule. Although probabilities are not identified with relative frequencies in subjective Bayesianism, it still offers an account of statistical inference wherein observing relative frequencies causes probabilities to be updated. If certain technical conditions hold, probability assignments will converge to the observed relative frequency in the limit of a large number of trials. Therefore, we could assert that noncontextuality is a brute empirical fact<sup>9</sup>.

The problem with this is that it provides no explanation of why noncontextuality holds. If we accept this, we might as well just give up and say that the only reason why we believe any physical theory is because it matches the observed relative frequencies. It would be like saying that the reason why the Maxwell-Boltzmann distribution applies to a box of gas is because we have sampled many molecules from such boxes and always found them to be approximately Maxwell-Boltzmann distributed. This belies the important explanatory role of stationary distributions in equilibrium statistical mechanics, and would be of no help in understanding why nonequilibrium systems tend to equilibrium. Similarly, the Born rule appears to be playing an important structural role in quantum theory that calls for an explanation.

The remaining option is to view noncontextuality as arising from physical, as opposed to logical, equivalence. The Dutch book rationality criterion is usually expressed as the requirement that you should not enter into bets that lead to a sure loss by logical necessity, but it is equally irrational to enter into a bet that you believe will lead to a sure loss, whether or not that belief is a logical necessity. Because of this, the argument that you should assign probability one to the certain event equally applies to events that you only

<sup>8</sup>Pitowsky attempts to argue along these lines [22], unsuccessfully in my view.

<sup>9</sup>This has been suggested in the context of the many-worlds interpretation [28].

believe to be certain, regardless of whether that belief is correct. Now, belief in the laws of physics entails certainty about statements that follow from the laws so this can be the origin of constraints on probability assignments.

To illustrate, suppose you believe that Newtonian mechanics is true and that there is a single particle system with a given Hamiltonian. This means that you are committed to propositions of the form “If the particle initially occupies phase space point  $(x_0, p_0)$  then at time  $t$  it occupies the solution to Hamilton’s equations  $(x(t), p(t))$  with initial condition  $(x_0, p_0)$ ”. If you bet on such propositions at anything more or less than even odds then you believe that you will lose money with certainty. Importantly, this type of argument can also imply constraints on events that you are not certain about. For example, if you assign a phase space region some probability and then compute the endpoints of the trajectories for all points in that region at a later time then the region formed by the end points must be assigned the same probability at that later time. This shows that the need to assign equal probabilities to different events can sometimes be derived from the laws of physics.

Crucially, this sort of argument can only really be made to work if there is an objectively existing external reality. There needs to be some sort of “quantum stuff” such that events that are always assigned the same probability correspond to physically equivalent states of this stuff. In the context of the many-worlds interpretation, the Deutsch-Wallace [18, 19, 20, 21] and Zurek [31, 32] derivations of the Born rule are arguments of this type, where the quantum stuff is simply the wavefunction.

## 7 “It from bit” or “bit from it”?

We have arrived at the conclusion that noncontextuality must be derived in terms of an analysis of the things that objectively exist. This implies a realist view of physics, or in other words “bit from it”, which seems to conflict with “it from bit”. Fortunately, this conflict is only apparent because “it” is being used in different senses in “it from bit” and “bit from it”. The things that Wheeler classifies as “it” are things like particles, fields and spacetime. They are things that appear in the fundamental ontology of classical physics and hence are things that only appear to be real from our perspective as classical agents. He does not mention things like wavefunctions, subquantum particles, or anything of that sort. Thus, there remains the possibility that reality is made of quantum stuff and that the interaction of this stuff with our question asking apparatus, also made of quantum stuff, is what causes the answers (particles, fields, etc.) to come into being. “It from bit” can be maintained in this picture provided the answers depend not only on the state of the system being measured, but also on the state of the stuff that comprises the measuring apparatus. Thus, we would end up with “it from bit from it”, where the first “it” refers to classical ontology and the second refers to quantum stuff.

## 8 Conclusion

On the subjective Bayesian view, “it from bit” implies that probability theory needs to be generalized, which is in accord with the observation that quantum theory is a generalized probability theory. However, “it from bit” does not get us all the way to the quantum probability rule. A subjective Bayesian analysis of noncontextuality indicates that it can only be derived within a realist approach to physics. At present, this type of derivation has only been carried out in the many-worlds interpretation, but I expect it can be made to work in other realist approaches to quantum theory, including those yet to be discovered.

## References

- [1] J. A. Wheeler, *Proceedings of the 3rd International Symposium on Foundations of Quantum Mechanics in the Light of New Technology*, S. Kobayashi, H. Ezawa, Y. Murayama, S. Nomura, eds. (Physical Society of Japan, Tokyo, 1990), pp. 354-368.

- [2] J. A. Wheeler, *Problems in the Foundations of Physics: Proceedings of the International School of Physics "Enrico Fermi," Course LXXII*, G. Toraldo di Francia, ed. (North-Holland, Amsterdam, 1979), pp. 395–492.
- [3] G. Birkhoff, J. von Neumann, *Annals of Mathematics* **37**, 823 (1936).
- [4] F. J. Murray, J. von Neumann, *Annals of Mathematics* **37**, 116 (1936).
- [5] M. Rédei, S. J. Summers, *Stud. Hist. Phil. Mod. Phys.* **38**, 390 (2007).
- [6] M. S. Leifer, R. W. Spekkens, Formulating quantum theory as a causally neutral theory of Bayesian inference (2011). Eprint arXiv:1107.5849.
- [7] D. Gillies, *Philosophical Theories of Probability* (Routledge, 2000).
- [8] A. Eagle, ed., *Philosophy of probability: contemporary readings* (Routledge, 2011).
- [9] B. de Finetti, *Fundamenta Mathematicae* **17**, 298 (1931).
- [10] F. P. Ramsey, *The Foundations of Mathematics and Other Logical Essays*, R. B. Braithwaite, ed. (Routledge and Kegan Paul, 1931), pp. 156–198. Reprinted in [8].
- [11] R. Jeffrey, *Subjective Probability: The Real Thing* (Cambridge University Press, 2004).
- [12] C. Howson, P. Urbach, *Scientific Reasoning: The Bayesian Approach* (Open Court, 2005), third edn.
- [13] D. Lewis, *Studies in Inductive Logic and Probability*, R. C. Jeffrey, ed. (University of California Press, 1980). Reprinted with postscript in [8].
- [14] C. M. Caves, C. A. Fuchs, R. Schack, *Phys. Rev. A* **65**, 022305 (2002).
- [15] C. A. Fuchs, *J. Mod. Opt.* **50**, 987 (2003).
- [16] C. A. Fuchs, *Physics in Canada* **66**, 77 (2010). Eprint arXiv:1003.5182.
- [17] C. A. Fuchs, QBism, the Perimeter of quantum Bayesianism (2010). Eprint arXiv:1003.5209.
- [18] D. Deutsch, *Proc. R. Soc. Lond. A* **455**, 3129 (1999). Eprint arXiv:quant-ph/9906015.
- [19] D. Wallace, *Stud. Hist. Phil. Mod. Phys.* **38**, 311 (2007). Eprint arXiv:quant-ph/0312157.
- [20] D. Wallace, *Many Worlds? Everett, Quantum Theory, and Reality*, S. Saunders, J. Barrett, A. Kent, D. Wallace, eds. (Oxford University Press, 2010). Eprint arXiv:0906.2718.
- [21] D. Wallace, *The Emergent Multiverse* (Oxford University Press, 2012).
- [22] I. Pitowsky, *Stud. Hist. Phil. Mod. Phys.* **34**, 395 (2003). Eprint arXiv:quant-ph/0208121.
- [23] B. de Finetti, *Philosophical Lectures on Probability*, vol. 340 of *Synthese Library* (Springer, 2008).
- [24] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, 2003).
- [25] D. Kahneman, *Thinking, Fast and Slow* (Penguin, 2012).
- [26] C. A. B. Smith, *J. Roy. Stat. Soc. B Met.* **23**, 1 (1961).
- [27] L. J. Savage, *The Foundations of Statistics* (Dover, 1972), second edn.
- [28] H. Greaves, W. Myrvold, *Many Worlds?: Everett, Quantum Theory, and Reality*, S. Saunders, J. Barrett, A. Kent, D. Wallace, eds. (Oxford University Press, 2010), chap. 9, pp. 264–306.
- [29] J. Barrett, *Phys. Rev. A* **75**, 032304 (2007). Eprint arXiv:quant-ph/0508211.
- [30] S. Kochen, E. P. Specker, *Journal of Mathematics and Mechanics* **17**, 59 (1967).
- [31] W. H. Zurek, *Phys. Rev. Lett.* **90**, 120404 (2003). Eprint arXiv:quant-ph/0211037.
- [32] W. H. Zurek, *Phys. Rev. A* **71**, 052105 (2005). Eprint arXiv:quant-ph/0405161.
- [33] A. M. Gleason, *Journal of Mathematics and Mechanics* **6**, 885 (1957).



## A Technical Endnotes

In general, a betting context  $\mathcal{B}$  is a Boolean algebra, which we take to be finite for simplicity. All such algebras are isomorphic to the algebra generated by the subsets of some finite set  $\Omega_{\mathcal{B}}$ , where AND is represented by set intersection, OR by union, and NOT by complement.

In quantum theory, a betting context corresponds to a set of measurements that can be performed together that is as large as possible. A measurement is represented by a self-adjoint operator  $M$  and all such operators have a spectral decomposition of the form

$$M = \sum_j \lambda_j \Pi_j, \quad (1)$$

with eigenvalues  $\lambda_j$  and orthogonal projection operators  $\Pi_j$  that sum to the identity  $\sum_j \Pi_j = I$ . The eigenvalues are the possible measurement outcomes and, when the system is assigned the density operator  $\rho$ , the Born rule states that the outcome  $\lambda_j$  is obtained with probability

$$p(\lambda_j) = \text{Tr}(\Pi_j \rho). \quad (2)$$

The eigenvalues just represent an arbitrary labelling of the measurement outcomes, so a measurement can alternatively be represented by a set of orthogonal projection operators  $\{\Pi_j\}$  that sum to the identity  $\sum_j \Pi_j = I$ , which is sometimes known as a *Projection Valued Measure (PVM)*<sup>10</sup>.

Two PVMs  $A = \{\Pi_j\}$  and  $B = \{\Pi'_k\}$  can be measured together if and only if each of the projectors commute, i.e.  $\Pi_j \Pi'_k = \Pi'_k \Pi_j$  for all  $j$  and  $k$ . If this is the case then  $\Pi_j \Pi'_k$  is also a projector and  $\sum_{jk} \Pi_j \Pi'_k = I$ . Therefore, one way of performing the joint measurement is to measure the PVM  $C = \{\Pi''_{jk}\}$  with projectors  $\Pi''_{jk} = \Pi_j \Pi'_k$  and, upon obtaining the outcome  $(jk)$ , report the outcome  $j$  for  $A$  and  $k$  for  $B$ . This fine graining procedure can be iterated by adding further commuting PVMs and forming the product of their elements with those of  $C$ . The procedure terminates when the resulting PVM is as fine grained as possible and this will happen when it consists of rank-1 projectors onto the elements of an orthonormal basis. The outcome of any other commuting PVM is determined by coarse graining the projectors onto the orthonormal basis elements.

Therefore, in quantum theory, we can take the sets  $\Omega_{\mathcal{B}}$  that generate the betting contexts  $\mathcal{B}$  to consist of the elements of orthonormal bases. An event  $E \in \mathcal{B}$  is then a subset of the basis elements and corresponds to a projection operator  $\Pi_E = \sum_{|\psi\rangle \in E} |\psi\rangle\langle\psi|$ . The Boolean operations on  $\mathcal{B}$  can be represented in terms of these projectors as

- Conjunction:  $G = E \text{ AND } F \Rightarrow \Pi_G = \Pi_E \Pi_F$
- Disjunction:  $G = E \text{ OR } F \Rightarrow \Pi_G = \Pi_E + \Pi_F - \Pi_E \Pi_F$ , which reduces to  $\Pi_G = \Pi_E + \Pi_F$  when  $E \cap F = \emptyset$ .
- Negation:  $G = \text{NOT } E \Rightarrow \Pi_G = I - \Pi_E$ .

From the Dutch book argument applied within a betting context, we have that our degrees of belief should be represented by a set of probability measures  $p(E|\mathcal{B})$  satisfying

- For any event  $E \subseteq \Omega_{\mathcal{B}}$ ,  $p(E|\mathcal{B}) \geq 0$ .
- For the certain events  $\Omega_{\mathcal{B}}$ ,  $p(\Omega_{\mathcal{B}}|\mathcal{B}) = 1$
- For disjoint events within the same betting context  $E, F \subseteq \Omega_{\mathcal{B}}$ ,  $E \cap F = \emptyset$ ,  $p(E \cup F|\mathcal{B}) = p(E|\mathcal{B}) + p(F|\mathcal{B})$ .

---

<sup>10</sup>More generally, we could work with Positive Operator Valued Measures (POVMs) or sets of consistent histories, but this would not substantially change the arguments of this essay.

The Born rule is an example of such an assignment, and in this language it takes the form

$$p(E|\mathcal{B}) = \text{Tr}(\Pi_E \rho). \quad (3)$$

The Born rule also has the property that the probability only depends on the projector associated with an event, and not on the betting context that it occurs in. For example, in a three dimensional Hilbert space, consider the betting contexts  $\Omega_{\mathcal{B}} = \{|0\rangle, |1\rangle, |2\rangle\}$  and  $\Omega_{\mathcal{B}'} = \{|+\rangle, |-\rangle, |2\rangle\}$ , where  $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$ . The Born rule implies that  $p(\{|2\rangle\}|\mathcal{B}) = p(\{|2\rangle\}|\mathcal{B}')$  and also that  $p(\{|0\rangle, |1\rangle\}|\mathcal{B}) = p(\{|+\rangle, |-\rangle\}|\mathcal{B}')$  because, in each case, the events correspond to the same projectors. The Dutch book argument alone does not imply this because it does not impose any constraints across different betting contexts.

A probability assignment is called *noncontextual* if  $p(E|\mathcal{B}) = p(F|\mathcal{B}')$  whenever  $\Pi_E = \Pi_F$ . Gleason's theorem [33] says that, in Hilbert spaces of dimension 3 or larger, noncontextual probability assignments are exactly those for which there exists a density operator  $\rho$  such that  $p(E|\mathcal{B}) = \text{Tr}(\Pi_E \rho)$ , i.e. they must take the form of the Born rule. Therefore, the Born rule follows from the conjunction of the Dutch book constraints and noncontextuality, at least in Hilbert spaces of dimension 3 or greater.