Bayesian Conditioning, the Reflection Principle, and Quantum Decoherence[*]

Christopher A. Fuchs and Rüdiger Schack

May 29, 2018

Abstract

The probabilities a Bayesian agent assigns to a set of events typically change with time, for instance when the agent updates them in the light of new data. In this paper we address the question of how an agent's probabilities at different times are constrained by Dutch-book coherence. We review and attempt to clarify the argument that, although an agent is not forced by coherence to use the usual Bayesian conditioning rule to update his probabilities, coherence does require the agent's probabilities to satisfy van Fraassen's [1984] *reflection principle* (which entails a related constraint pointed out by Goldstein [1983]). We then exhibit the specialized assumption needed to recover Bayesian conditioning from an analogous reflection-style consideration. Bringing the argument to the context of quantum measurement theory, we show that "quantum decoherence" can be understood in purely personalist terms—quantum decoherence (as supposed in a von Neumann chain) is not a physical process at all, but an application of the reflection principle. From this point of view, the decoherence theory of Zeh, Zurek, and others as a story of quantum measurement has the plot turned exactly backward.

1 Introduction

At the center of most accounts of Bayesian probability theory [4] is the procedure of Bayesian conditioning. By this we mean the following. Assume a Bayesian agent, at some time $t = 0$, has assigned probabilities $P_0(E)$, $P_0(D)$ and $P_0(E, D)$ to events E and D and their conjunction. As long as $P_0(D) \neq 0$, the conditional probability of E given D is then

$$P_0(E|D) = \frac{P_0(E, D)}{P_0(D)} \ . \tag{1}$$

Now assume that, at a later time $t = \tau$, the agent learns that D is true and updates his probability for E. We denote the agent's updated probability by $P_\tau(E)$. Standard

---

[*]This article is dedicated to the memory of Itamar Pitowsky. Already 10 years ago, he saw the potential in a Bayesian understanding of quantum probabilities and had a respect for our efforts like no one else [1, 2, 3]. The respect was mutual. For you, Itamar.

Bayesian conditioning consists of setting

$$P_\tau(E) = P_0(E|D) . \tag{2}$$

The rule Eq. (2) can be viewed as a possible answer to the general question of how an agent's probabilities at two different times should be related. We will address this question from a personalist Bayesian perspective [4, 5, 6, 7, 8, 9], according to which probabilities express an agent's uncertainty, or degrees of belief, about future events and acquire an operational meaning through decision theory [4]. Although they are not determined by agent-independent facts, personalist probability assignments are not arbitrary. Dutch-book coherence [5, 8, 9] as a normative principle requires an agent to try her best to make her numerical belief assignments conform to the usual rules of the probability calculus. When coupled with the agent's overall belief system, this is a powerful constraint [10]. Personalist Bayesian probability is at the heart of Quantum Bayesianism, a radical new approach to the foundations of quantum mechanics developed by Caves, Fuchs, Schack, Appleby, Barnum, and others. (See [11, 12] for an extensive reference list.) The motivation for the present investigation is to explore the relevance of Bayesian conditioning to the Quantum Bayesian program.

It was first pointed out by Hacking [13] that there is no coherence argument that compels the agent to take into account the earlier probabilities $P_0(E|D)$ when setting the later probabilities $P_\tau(E)$. Similar points have been made by other authors (see, e.g., [14]). Hacking was writing about the standard *synchronic* Dutch book arguments, but the above statement remains true even for the *diachronic* Dutch book arguments, originally due to Lewis and first reported by Teller [15]. Without further assumptions, diachronic coherence does not compel the agent, at $t = \tau$, to use the Bayesian conditioning rule (1).

The way diachronic coherence arguments connect probability assignments at different times is more subtle. It is expressed elegantly through van Fraassen's reflection principle [16], which itself entails the related constraints of Shafer [17] and Goldstein [18]. The key idea behind the reflection principle is to consider the agent's beliefs about his own future probabilities, i.e., to consider expressions such as $P_0(P_\tau(E) = q)$. Shafer [17] put the point very nicely,

> This interpretation is based on the assumption that a person has subjective probabilities for how his information and probabilities may change over time. This means we are concerned not with how the person *should* or *will* change his beliefs, but rather with what he believes about how these beliefs will change. [Emphases added.]

The same idea underlies the approach this paper takes towards Bayesian conditioning and quantum decoherence.

In the next section, we present a detailed example where the agent appears justified to depart from the Bayesian conditioning rule. In Section 3 we review the standard, synchronic, Dutch book arguments and show why they do not imply Bayesian conditioning. Section 4 introduces diachronic coherence and presents a derivation of the reflection principle. In Section 5 we show that the Bayesian conditioning rule can be understood as a variant of the reflection principle valid for a particular class of situations. Section 6 addresses an argument that has been advanced against the reflection principle

and shows that it is based on a misconception of the role that coherence considerations have in probability assignments. Finally, in Section 7 we give a natural application of the reflection principle to decoherence in quantum mechanics from a Quantum Bayesian perspective.

2 Example: Polarization Data

Consider a physicist running an experiment to discover the linear polarization of photons coming from a rather complicated optical device which he had built himself. Perhaps he is convinced that every photon is produced the same way, only that he has forgotten which orientation θ he gave to a certain polarization filter deep within the set-up. It might thus be easier to discover θ and recalibrate than to tear the whole thing apart and readjust. A statistical analysis is in order.

Our experimenter will measure polarization for a sequence of n individual photons and carry out a Bayesian analysis of the outcomes consisting of a string, s_n, of zeros and ones. The zeros stand for horizontal polarization and ones for vertical polarization. Before starting, the experimenter records his probabilistic prior, which in the personalist approach to probability adopted in this paper, represents his Bayesian degrees of belief about the measurement outcomes. To be specific we assume that the prior is *exchangeable*, i.e., of the form [8, 19]

$$P_0(s_n) = \int_0^1 q_0(x) x^k (1-x)^{n-k} dx \ , \tag{3}$$

where k is the number of zeros in s_n, and $q_0(x)$ is a probability density. If the experimenter is completely ignorant of what orientation he had given the filter, he might assume $q_0(x)$ to be the constant density, but the precise form of $q_0(x)$ is of no great importance to the argument below.

The prior $P_0(s_n)$ is a convex sum of binomial distributions $\{x, 1-x\}$, with $x = \cos^2 \theta$. It is symmetric in the sense that it is invariant under permutations of the bits in s_n. By adopting this prior, the experimenter judges that the order in which the photons arrive is irrelevant to his analysis. To a Bayesian, this is in fact the operational meaning that all photons are "produced the same way." A simple consequence of this is that a posterior probability calculated from this prior by Bayesian conditionalization after some number of trials will not depend on the order of the zeros and ones found in those trials.

Suppose now that the experimenter observes 4000 trials, and finds very nicely that the number of zeros and ones is not very far from 2000 each. Technically this means that if the experimenter updates his prior to a posterior by Bayesian conditionalization, the posterior for the next n bits will be

$$P_{\text{posterior}}(s_n) = \int_0^1 q'(x) x^k (1-x)^{n-k} dx \ , \tag{4}$$

where $q'(x)$ is a function on the interval $[0, 1]$ peaked very near $1/2$. If the experimenter were asked to bet on the next bit, this probability distribution would advise him to bet at close to even odds.

But now consider the following fantastic scenario. Suppose the experimenter becomes aware that the string s_n he accumulated is identical to the first 4000 bits of the binary expansion of π! Any sane person would be flabbergasted. Even though the experimenter built the device with his own hands, he would surely wonder what was up. Perhaps one of his lab partners has played an immense joke on him?

The following question becomes immediate: If the experimenter is rational, how should he bet *now* on the next bit? Sticking doggedly with Bayesian conditioning, he would be advised to use near even odds, just as before. But the number π is too significant to ignore: His heart says to bet with

$$P_\tau(E) = 0.99 \tag{5}$$

on the event E that the next bit equals the 4001th bit of the binary expansion of π. Our experimenter faces a stark choice: He can either ignore his heartfelt belief and use the value $P_\tau(E) \approx 1/2$ obtained from the conditioning rule, or he can ignore the conditioning rule and use the value $P_\tau(E) = 0.99$ representing what he really feels. Both choices are deeply problematical: the second one seems to be incoherent because it is contradicted by the usual understanding of the formalism, whereas the first one seems to be ignoring common sense to the point of being foolish.

Is the second choice really incoherent however? Does it violate a reasonable normative requirement? We will see in the next sections that this is not the case.

3   Synchronic Dutch Book

A simple way to give an operational meaning to personalist Bayesian probabilities is through the agent's betting preferences. When an agent assigns probability p to an event E, she regards \$$p$ to be the fair price of a standard lottery ticket that pays \$1 if E is true. In other words, an agent who assigns probability p to the event E regards both buying and selling a standard lottery ticket for \$$p$ as fair transactions; for her, the ticket is worth \$$p$.

In the following, we will call a set of probability assignments *incoherent* if it can lead to a sure loss in the following sense: there exists a combination of transactions consisting of buying and/or selling a finite number of lottery tickets which (i) lead to a sure loss and (ii) the agent regards as fair according to these probability assignments. A set of probability assignments is *coherent* if it is not incoherent. We accept as a normative principle that an agent should aim to avoid incoherent probability assignments.

The standard, synchronic, Dutch book argument [5, 8, 9] shows that an agent's probability assignments P_0 at a given time are coherent if and only if they obey the usual probability rules, i.e., $0 \le P_0(E) \le 1$ for any event E; $P_0(S) = 1$ if the agent believes the event S to be true; and $P_0(E \vee D) = P_0(E) + P_0(D)$ for any two events E and D that the agent believes to be mutually exclusive.

In this approach, conditional probability is introduced as the fair price of a lottery ticket that is refunded if the condition turns out to be false. Formally, let D and E be events, and let \$$q$ be the price of a lottery ticket that pays \$1 if both D and E are true, and \$$q$ (thus refunding the original price) if D is false. For the agent to make the

conditional probability assignment $P_0(E|D) = q$ means that she regards \$$q$ to be the fair price of this ticket.

It is then a consequence of Dutch-book coherence that the product rule $P_0(E, D) = P_0(E|D)P_0(D)$ must hold [9]. In other words, conditional probability assignments violating this rule are incoherent. If $P_0(D) \neq 0$, we obtain Bayes's rule,

$$P_0(E|D) = \frac{P_0(E, D)}{P_0(D)} \ . \tag{6}$$

It is worth pointing out that Bayes's rule emerges here as a theorem, combining terms that are defined independently, in contrast to the common axiomatic approach to probability theory where Eq. (6) is used as the definition of conditional probability.

The above shows that a coherent agent must use Bayes's rule to set the conditional probability $P_0(E|D)$. The value of $P_0(E|D)$ expresses what ticket prices the agent regards as fair at time $t = 0$, i.e., before she finds out the truth value of either D or E. It says nothing about what ticket prices she will regard as fair at some later time $t > 0$.

In particular, assume that, at some time $t = \tau > 0$, the agent learns that D is true and updates her probabilities accordingly. Denote by $P_\tau(E)$ the agent's updated probability of E, meaning that she now regards \$$P_\tau(E)$ as the new fair price of a ticket that pays \$1 if and only if E is true. Nothing in the Dutch book argument sketched above implies that $P_\tau(E)$ should be equal to $P_0(E|D)$ [13]. All probabilities used in the argument are the agent's probabilities at time $t = 0$; they are defined via ticket prices for bets on E, D and their conjunction which the agent regards as fair at $t = 0$. The Dutch book argument leading to Eq. (6) is a *synchronic* argument. It does not connect in any way the agent's probability assignments at $t = 0$ and $t = \tau$. In particular, it does not imply that the agent has to use Bayesian conditioning to update her probabilities. In the next section we will see what connection between the agent's probability assignments at different times actually is implied by diachronic Dutch book arguments.

4   Diachronic Dutch Book

To set the scene, we consider an investor who today buys 500 shares of some company at a price of \$20 each, which he regards as a fair deal. The next day, his appreciation of the market has changed, and he sells his 500 shares at \$18 each, which now, given the new situation, he again regards as a fair deal. Despite the fact that the investor makes a net loss of \$1000, he does not behave irrationally. By selling his shares at a lower price on the next day, he simply cuts his losses.

But what if the investor is certain today that tomorrow he will regard \$18 as the fair price for a share? It would then be foolish for him to buy, today, 500 shares for \$20 each, because he is certain that tomorrow he would be willing to sell the shares for \$18 each, leading to a net loss of \$1000. As a matter of fact, buying shares at any price above \$18 today would be foolish in this situation, as would be selling shares today at any price below \$18.

In the above example, we have assumed that money has the same utility for the investor today and tomorrow, i.e., we have assumed a zero interest rate. This is an

assumption we will make throughout the present paper. More precisely, we will assume that the time at which she receives a sum of money is irrelevant to a Bayesian agent.

In probability language, what we have just described is the following. Assume $P_0(E)$ is an agent's probability at $t = 0$ of some event E. The agent buys a lottery ticket that pays \$1 if E is true, for \$$P_0(E)$ which she regards as the fair price. At a later time $t = \tau$, she updates her beliefs. Her probability for E is now $P_\tau(E)$, which happens to be less than $P_0(E)$. At this point, the agent decides to cut her losses by selling the ticket for \$$P_\tau(E)$. Despite the net loss, there is nothing irrational about the agent's transactions.

But now suppose that at $t = 0$ the agent is certain that, at $t = \tau$, her probability of E will be q, where $q \neq P_0(E)$. In the case $q < P_0(E)$, this means that, at $t = 0$, she is willing to buy a ticket for \$$P_0(E)$ although she already knows that later she will be willing to sell it for the lower price \$$q$. In the case $q > P_0(E)$, it means that, at $t = 0$, she is willing to sell a ticket for \$$P_0(E)$ although she already knows that later she will be willing to buy the same ticket for the higher price \$$q$. In both cases, already at $t = 0$ the agent is certain of a sure loss.

This simple scenario contains the main idea of van Fraassen's diachronic Dutch book argument. Similar to the synchronic case discussed in the previous section, we call an agent's probability assignments incoherent if there exists a combination of transactions consisting of buying and/or selling a finite number of lottery tickets at two different times such that (i) *already at the earlier time*, the agent is sure of a net loss; and (ii) each transaction is regarded as fair by the agent according to her probability assignments *at the time* the transaction takes place. We will continue to accept as a normative principle that an agent should aim to avoid incoherent probability assignments.

To turn our simple scenario into the full-fledged diachronic Dutch-book argument, we only need to relax the assumption that at $t = 0$ the agent is *certain* that $P_\tau(E) = q$. Instead, we assume that

$$P_0\big(P_\tau(E) = q\big) > 0 \ , \tag{7}$$

i.e, at $t = 0$ the agent believes with some positive probability that at $t = \tau$ her probability of E will be equal to q. We will now show that this implies the agent's probability assignments are incoherent unless

$$P_0\big(E \mid P_\tau(E) = q\big) = q \ , \tag{8}$$

i.e., unless at $t = 0$ the agent's conditional probability of E, given that $P_\tau(E) = q$, equals q. This is van Fraassen's *reflection principle* [16].

To derive the reflection principle, denote by Q the proposition $P_\tau(E) = q$, i.e., the assertion that at time $t = \tau$, the agent will regard \$$q$ as the fair price for a ticket that pays \$1 if and only if E is true. The inequality (7) thus becomes $P_0(Q) > 0$. To establish that coherence implies the reflection principle (8), one must show that the assumption $P_0(E|Q) \neq q$ leads to a sure loss for an appropriately chosen set of bets no matter what outcomes occur for the events considered.

As a warm-up to gain intuition, suppose that $P_0(E|Q) > q$ and that Q is true. This means that at $t = 0$, the agent is willing to buy a ticket for \$$P_0(E|Q)$ that pays \$1 if both Q and E are true, and refunds the ticket price if Q is false. But, because of Q's truth, this ticket will further be equivalent to a ticket that pays \$1 if E is true. Finally,

the truth of Q also implies that at $t = \tau$ the agent will be willing to sell this ticket for $\$q$, which is less than what she paid for it. In other words, if Q is true, the agent is sure to lose $\$d$, where $d = P_0(E|Q) - q$.

But this simple argument—illustrative though it may be—is not a full-fledged proof of incoherence. To get a full proof we need to show that the agent is sure of a loss not only when Q turns out to be true, but also when Q turns out to be false. For this it is sufficient to consider an alternate scenario where there is a side bet on Q, such that the agent loses some amount if Q is false, and wins less than $\$d$ if Q is true. Such a side bet may be realized by a lottery ticket that pays $\$d/2$ if Q is true, which the agent is willing to buy for $\$P_0(Q)d/2$.

We have thus the following combination of transactions, each of which the agent regards as fair at the time it takes place:

(i) to buy, at $t = 0$, for $\$P_0(E|Q)$, a ticket that pays $\$1$ if both Q and E are true, and refunds the ticket price if Q is false;

(ii) to buy, at $t = 0$, for $\$P_0(Q)d/2$, a lottery ticket that pays $\$d/2$ if Q is true;

(iii) if Q is true (i.e., if $P_\tau(E) = q$), to sell, at $t = \tau$, for $\$q$, a lottery ticket that pays $\$1$ if E is true.

Already at $t = 0$, the agent knows that these transactions result in a net loss, equal to $\$(P_0(Q) + 1)d/2$ if Q is true, and $\$dP_0(Q)/2$ if Q is false. We have thus shown that the assumption $P_0(E|Q) > q$ implies that the agent's probability assignments are incoherent.

The final piece of a proof is to consider the case $P_0(E|Q) < q$. By reversing the signs of all transactions above, it is easy to see that this case leads to a sure loss in exactly the same way. Putting these two cases together, this completes the full derivation of the reflection principle.

The coherence condition of Shafer [17] and Goldstein [18] follows by a simple application of synchronic coherence along with the reflection principle. Suppose the agent instead of contemplating a single $Q = [P_\tau(E) = q]$ for what she will believe of E at $t = \tau$, contemplates a range of mutually exclusive and exhaustive propositions $\{Q\}$ to which she assigns probabilities $P_0(Q)$. Then, straightforward synchronic coherence requires

$$P_0(E) = \sum_Q P_0(Q)P_0(E|Q) = \sum_q P_0\big(P_\tau(E) = q\big)\, P_0\big(E \mid P_\tau(E) = q\big)\,, \qquad (9)$$

for which reflection in turn implies

$$P_0(E) = \sum_q P_0\big(P_\tau(E) = q\big)\, q\,. \qquad (10)$$

This implication of the reflection principle will turn out to be particularly important for our exposition of quantum decoherence.

5   Bayesian Conditioning in Reflectional Terms

The reflection principle is a constraint on an agent's present beliefs about her future probability assignments. It does not directly provide an explicit rule for assigning probabilities, either for the present ones or the future ones. An agent whose probabilities

violate the reflection principle is incoherent and should strive to remove this incoherence. The reflection principle does not provide a recipe for how to do this.

One way in which the agent can achieve coherence is by adopting an "updating strategy" [20] based on the Bayesian conditioning rule. We will now explore to what extent the Bayesian conditioning rule follows for such a strategy in a way analogous to the reflection principle—that is, in a way "concerned not with how the person should or will change his beliefs, but rather with what he believes about how these beliefs will change" [17].

Let E and D be events, and let $P_0(E)$, $P_0(D)$ and $P_0(E|D)$ denote the agent's respective probabilities at $t = 0$. Assume that the truth value of D will be revealed to the agent at $t = \tau$. Suppose she now adopts the strategy that, if at $t = \tau$ she learns that D is true, her updated probability of E, denoted by $P_\tau(E)$, will be given by some value q, $0 \leq q \leq 1$.

The above can be phrased in terms of the agent's probabilities at $t = 0$. For her to adopt this strategy simply means that she is certain that, if D turns out to be true, she will make the probability assignment $P_\tau(E) = q$, i.e.,

$$P_0\big(P_\tau(E) = q \big| D\big) = 1 . \tag{11}$$

This statement about the agent's current belief about her future probability captures the essence of Bayesian conditioning. Together with diachronic coherence it implies that

$$q = P_0(E|D) , \tag{12}$$

i.e., the Bayesian conditioning rule. Presented in this way, it can be regarded as a variant of the reflection principle, valid whenever the condition (11) holds.

To derive Eq. (12), we consider again the combinations of bets introduced in Sec. 4, but with the event D replacing Q throughout. We assume first that $P_0(E|D) > q$ and define $d = P_0(E|D) - q$. The transactions are

(i) to buy, at $t = 0$, for $\$P_0(E|D)$, a ticket that pays $\$1$ if both D and E are true, and refunds the ticket price if D is false;

(ii) to buy, at $t = 0$, for $\$P_0(D)d/2$, a lottery ticket that pays $\$d/2$ if D is true;

(iii) if D is true, to sell, at $t = \tau$, for $\$q$, a lottery ticket that pays $\$1$ if E is true.

At $t = 0$, the agent is certain that these transactions result in a net loss, equal to $\$(P_0(D) + 1)d/2$ if D is true, and $\$P_0(D)d/2$ if D is false. At $t = 0$, the agent regards (i) and (ii) as fair transactions, and because of Eq. (11), she is certain that at $t = \tau$ she will regard (iii) as a fair transaction. We have thus shown that the agent's probabilities are incoherent if $P_0(E|D) > q$. The case $P_0(E|D) < q$ is similar. Thus coherence implies that $P_0(E|D) = q$, as required.

The key assumption in this derivation, expressed by Eq. (11), is that the agent can identify an event D that she expects to determine her future beliefs. There are more general updating strategies that are not of this form. Jeffrey's probability kinematics [21] is such an example. Probability kinematics is a coherent updating strategy [20] which

does not make use of the Bayesian conditioning rule, but it too can be put in reflectional terms as we did with Bayesian conditioning.

Actually, one can go still further along these reflectional lines if one strengthens the assumption in Eq. (11) to also make a direct identification between the possible values for $P_\tau(E)$ and the D, i.e., that there is bijection between them. In such a case, one can say that Bayesian conditioning *follows directly* from the reflection principle. For then,

$$P_0(E|D) = P_0(E|D,Q) = P_0(E|Q) \tag{13}$$

by standard synchronic logic, and $P_0(E|Q) = q$ by reflection.

The discussion above is entirely in terms of the agent's beliefs at $t = 0$. What if, at $t = \tau$, after learning that D is true, the agent re-analyses the situation, possibly taking into account circumstances she was not aware of at $t = 0$, and concludes that her new probability, $P_\tau(E)$, differs from $P_0(E|D)$. Does this imply that her probability assignment is incoherent? The answer is no. Coherence is a condition about an agent's *current beliefs*, including her beliefs about her future probability assignments. In the above scenario, the agent's beliefs at $t = \tau$ are coherent as long as $0 \leq P_\tau(E) \leq 1$. Nothing in the Dutch book argument implies that the agent's actual probabilities at $t = \tau$ are constrained by her probabilities at $t = 0$, which supports the conclusion of Sec. 2 that there is no conflict with coherence for an experimenter who assigns $P_\tau(E) = 0.99$ although the Bayesian conditioning rule appears to mandate $P_\tau(E) \approx 1/2$.

6   Sirens, Car Keys, and Married Couples

We have seen in the previous section that one way of satisfying the reflection principle and thereby avoiding incoherence is to set your future probabilities in terms of your current probabilities via Bayesian conditioning. The form of the reflection principle, however, suggests a different way of proceeding. Since Eq. (8) expresses a constraint on a current probability, conditioned on a future probability assignment, one could take the future probability as given and set the current probability in terms of it, thus reversing the usual direction of Bayesian updating. This can be a useful and legitimate procedure. An important application will be given in Sec. 7 below.

If the reflection principle is taken as a rule to set future probabilities in terms of current probabilities, it can lead to decisions that appear irrational [22, 23, 24, 25]. A classic example [26] is provided by the story of Ulysses and the Sirens. Ulysses knows that tomorrow, as soon as he is within earshot of the Sirens, he will make a catastrophic decision. Does the reflection principle force him to endorse this catastrophic decision today?

Since analysing this story would involve a discussion of utility, here is another famous example [23]. I know that I will get drunk this evening and that I will assign probability 10^{-6} to the event E of my causing an accident while driving home late at night. Does reflection imply that I must assign probability 10^{-6} to the event E now?

Examples like this have led, e.g., Christensen [23] to the conclusion that the reflection principle is unsound. This conclusion stems from a confusion about the role of coherence

arguments, however. The reflection principle can be regarded as a tool to detect incoherence. The Dutch book arguments show that incoherent probability assignments have the potential to lead to catastrophic consequences. This justifies accepting a normative rule that an agent should adhere to the reflection principle in order to avoid incoherence. The reflection principle does not, however, give a prescription for setting probabilities, either today's in terms of tomorrow's or vice versa. There is a range of options for the agent once she has detected an incoherence, as we will now illustrate.

Suppose that, in the example above, my initial conditional probability for an accident if I drive home drunk is 0.01. Suppose further that I am certain that I will get drunk, and that my probability for an accident will then be 10^{-6}. These probabilities violate the reflection principle. My probability assignments are therefore incoherent. One way of avoiding this incoherence would be to decide not to get drunk, which would mean assigning probability 0 to this event and therefore restore coherence. There is another very practical solution, which is to give my car keys to a trusted friend before I start drinking. My probability assignments will still be incoherent, but I will be unable to act on them.

Ulysses's solution, 3000 years ago, was very similar. He ordered his men to chain him to the mast of his ship. His men were to plug their ears. He accepted incoherence, but prevented himself from acting on his incoherent probability assignments. His men achieved coherence by reducing the probability of hearing the Sirens to zero. Coherence is an ideal one should always strive for. Incoherent probability assignments have the potential to lead to catastrophic consequences. If one can't achieve coherence, one should give up the car keys, plug one's ears or chain oneself to a mast.

In his article contra reflection, Christensen [23] pointed out that the reflection principle is very similar to a related principle which he called *solidarity*. Consider husband and wife who share a bank account. Denote the husband's probabilities by P_h and the wife's by P_w, and consider some event, E. Solidarity is the principle that, given that the wife's probability of E is q, the husband's probability must also be q, i.e.,

$$P_h\big(E\big|P_w(E) = q\big) = q \ . \tag{14}$$

Violating solidarity leads to a sure loss for the joint bank account exactly like in the diachronic Dutch book argument for the reflection principle.

Christensen argued that the solidarity principle is clearly absurd. This may be a case of confusion between normative and descriptive rules. The solidarity principle is a normative principle and does not claim that actual agents' probability assignments are always compatible with it. What it says instead is that, to avoid potential catastrophic consequences for their common bank account, husband and wife must strive for coherence. The solidarity principle, or more generally, the reflection principle, provides a tool to detect incoherence. It is then up to the agents how to resolve the incoherence. The husband might give, the wife might give, or they might compromise after debating all the relevant issues. The key point is that deliberation is to their mutual benefit, and coherence is their goal.

7   Quantum Decoherence

In the last section, we described how the reflection principle can be used to detect incoherence and thus to avoid catastrophic consequences. In this section, we will see that there is a generic situation in quantum theory where the reflection principle is used directly to set today's probabilities in terms of tomorrow's.

We will look at a standard quantum measurement situation [27] from the perspective of Quantum Bayesianism, according to which all quantum states, pure or mixed, represent an agent's degrees of belief about future measurement outcomes. Assume an agent has, at time $t = 0$, assigned a quantum state (i.e., density operator) ρ_0 to a quantum system. She intends to perform two measurements on the system, the first one at time $t = \tau > 0$, the second one at a still later time $t = \tau'$. She describes the first measurement by a collection of trace-decreasing completely positive maps $\{\mathcal{F}_i\}$, each corresponding to a potential outcome, i, for the first measurement. These completely positive maps determine the agent's probabilities $P_0(i) = \mathrm{tr}[\mathcal{F}_i(\rho_0)]$, at time $t = 0$, for the outcomes i, but they also determine the states she will assign to the system after the measurement: If outcome i, then $\rho_\tau = P_0(i)^{-1}\mathcal{F}_i(\rho_0)$.

To describe the second measurement, it is enough to use a POVM, i.e., positive operator valued measure, $\{E_j\}$, since we will not be considering any further measurements after it. In this description each positive operator E_j corresponds to a potential outcome, j, for the second measurement. If ρ_τ is the agent's system state at time $t = \tau$, then her probabilities, at $t = \tau$, for the outcomes j are given by $P_\tau(j) = \mathrm{tr}(E_j\rho_\tau)$.

Now suppose our agent is confronted at time $t = 0$ with a bet concerning the outcome j at $t = \tau'$. How should she gamble without having yet performed the measurement at $t = \tau$? We can read the answer straight off the reflection principle as written in the form of Goldstein and Shafer, Eq. (10)—remember here that $P_\tau(j)$ is implicitly dependent upon i:

$$P_0(j) = \sum_i P_0(i)P_\tau(j) = \sum_i \mathrm{tr}[E_j\mathcal{F}_i(\rho_0)] \,. \tag{15}$$

Cleaning this up a bit, we can write:

$$P_0(j) = \mathrm{tr}[E_j \sum_i \mathcal{F}_i(\rho_0)] = \mathrm{tr}(E_j\rho_0') \,, \tag{16}$$

where

$$\rho_0' = \sum_i \mathcal{F}_i(\rho_0) \,. \tag{17}$$

What we have shown here is that the reflection principle entails that the agent can obtain her probabilities at $t = 0$ for the outcomes of the second measurement from the density operator ρ_0'. The state ρ_0', which has the form of a "decohered" state, *is* the agent's quantum state at $t = 0$ *as far as the second measurement is concerned.*

These conclusions are valid for any pair of measurements, but a little more can be said if the POVM $\{E_j\}$ is informationally complete, i.e., if the state ρ_τ is fully determined by the probabilities $P_\tau(j)$. In this case ρ_0' as defined in Eq. (17) is the only density operator that gives rise to the probabilities $P_0(j)$ required by the reflection principle.

Equation (17) takes a perhaps more familiar form if the first measurement is a von Neumann measurement and the updating is given by the Lüders rule. In this case the

action of the maps $\mathcal{F}_i$ on the state ρ_0 can be written as $\mathcal{F}_i(\rho_0) = \Pi_i \rho_0 \Pi_i$, where the Π_i are projection operators, and Eq. (17) becomes

$$\rho_0' = \sum_i \Pi_i \rho_0 \Pi_i \ . \tag{18}$$

A common attitude about quantum measurement is that it is something that demands a detailed physical explanation. Much of the folklore since the publication of John von Neumann's 1932 book *Mathematical Foundations of Quantum Mechanics* is that a quantum measurement is something that occurs in two steps: First, there is a kind of "pre-measurement" where the quantum system becomes entangled with a measuring device. Secondly, there is a "selection" of one of the entangled state's components; this is what singles out a particular measurement result.

The trouble with this description, however, is that the entangled wave function, with its freedom to be expressed in any bipartite basis, does not have enough structure to specify how it should be decomposed so that such a "selection" can be effected. The theory of quantum decoherence, developed by Zeh, Zurek, and others [28], attempts to overcome this deficiency in the von Neumann story by supplementing it with a further story of interaction between the measuring device and an environment: The idea is that the specific form of the interaction with the environment specifies how the joint state of system plus device ought to be decomposed. In this picture, the decoherence process preceding the "selection" step leads to a state of the form (18), or more generally, (17). What remains mysterious in this picture, however, is "the selection step" itself. Decoherence theorists usually leave that question aside, implicitly endorsing one variety or another of an Everettian interpretation of quantum mechanics.

In contrast, the Quantum Bayesian view of quantum theory leaves most of the usual von Neumann story aside: Instead of taking quantum states and unitary evolution as the ontic elements to which the theory refers, it takes the idea of an individual agent's decisions and experience as the theory's real subject matter. In this view, the process called "quantum measurement" is nothing other than an agent acting upon the world and experiencing the consequences of her actions. For a Quantum Bayesian, the only *physical* process in a quantum measurement is what was previously seen as "the selection step"—i.e., the agent's action on the external world and its unpredictable consequence for her, the data that leads to a new state of belief about the system.

Thus, it would seem there is no foundational place for decoherence in the Quantum Bayesian program. And this is true. Nonetheless, in the two-time measurement scenario we described above, there is a coherent state assignment at time $t = 0$ for the second measurement that mimics a belief in decoherence. This is simply a consequence of the implications of the reflection principle. The "decohered" state ρ_0' is not the agent's state after she has made the first measurement (that would have been *one* of the ρ_τ depending upon the i found). It is not the state resulting from the measurement interaction before the selection step takes place as the decoherence program would have it (nothing is so intricately modeled here). It is simply a quantum state the agent uses at time $t = 0$ before the first measurement to make decisions regarding the outcomes of the second measurement.

That is the story of decoherence from the Quantum Bayesian perspective. Decoherence does not come conceptually before a "selection," but rather is predicated on a time

$t = 0$ belief regarding the possibilities for the next quantum state at time $t = \tau$. Decoherence comes conceptually *after* the recognition of the future possibilities. In this sense the decoherence program of Zeh and Zurek [28], regarded as an attempt to contribute to our understanding of quantum measurement, has the story exactly backward.

8   Acknowledgements

References

[1] I. Pitowsky, "Betting on the Outcomes of Measurements: A Bayesian Theory of Quantum Probability," Stud. His. Phil. Mod. Phys. **34**, 395 (2003).

[2] I. Pitowsky, "Quantum Mechanics as a Theory of Probability," in *Physical Theory and Its Interpretation: Essays in Honor of Jeffrey Bub*, edited by W. Demopoulos and I. Pitowsky, (Springer, Berlin, 2006), pp. 213–240.

[3] J. Bub and I. Pitowsky, "Two Dogmas about Quantum Mechanics," in *Many Worlds? Everett, Quantum Theory, and Reality*, edited by S. Saunders, J. Barrett, A. Kent, and D. Wallace, (Oxford University Press, Oxford, 2010), pp. 433–459.

[4] J. M. Bernardo and A. F. M. Smith, *Bayesian Theory*, (Wiley, Chichester, 1994).

[5] F. P. Ramsey "Truth and Probability," in F. P. Ramsey, *The Foundations of Mathematics and other Logical Essays*, edited by R. B. Braithwaite, (Harcourt, Brace and Company, New York, 1931), pp. 156–198.

[6] B. de Finetti, "Probabilismo," Logos **14**, 163 (1931); transl., "Probabilism," Erkenntnis **31**, 169 (1989).

[7] L. J. Savage, *The Foundations of Statistics*, (John Wiley & Sons, New York, 1954).

[8] B. de Finetti, *Theory of Probability*, 2 volumes, (Wiley, New York, 1990).

[9] R. Jeffrey, *Subjective Probability. The Real Thing*, (Cambridge University Press, Cambridge, 2004).

[10] J. Logue, *Projective Probability*, (Clarendon Press, Oxford, 1995).

[11] C. A. Fuchs and R. Schack, "Quantum-Bayesian Coherence," `arXiv:0906.2187v1`, (2009).

[12] C. A. Fuchs, "QBism, the Perimeter of Quantum Bayesianism," `arXiv:1003.5209v1`, (2010).

[13] I. Hacking, "Slightly More Realistic Personal Probability," Phil. Sci. **34**, 311 (1967).

[14] C. Howson and P. Urbach, *Scientific Reasoning: The Bayesian Approach*, second edition, (Open Court Publishing, La Salle, 1993).

[15] P. Teller, "Conditionalization and Observation," Synthese **26**, 218 (1973).

[16] B. C. van Fraassen, "Belief and the Will," J. Phil. **81**, 235 (1984).

[17] G. Shafer, "A Subjective Interpretation of Conditional Probability," J. Phil. Log. **12**, 453 (1983).

[18] M. Goldstein, "The Prevision of a Prevision," J. Am. Stat. Assoc. **78**, 817 (1983).

[19] C. M. Caves, C. A. Fuchs and R. Schack, "Unknown Quantum States: The Quantum de Finetti Representation," J. Math. Phys. **43**, 4537 (2002).

[20] B. Skyrms, "Dynamic Coherence and Probability Kinematics," Phil. Sci. **54**, 1 (1987).

[21] R. Jeffrey, *The Logic of Decision*, (McGraw Hill, New York, 1965).

[22] P. Maher, "Diachronic Rationality," Phil. Sci. **59**, 120 (1992).

[23] D. Christensen, "Clever Bookies and Coherent Beliefs," Phil. Rev. **100**, 229 (1991).

[24] B. Skyrms, "A Mistake in Dynamic Coherence Arguments?," Phil. Sci. **60** 320 (1993).

[25] M. S. Green and C. R. Hitchcock, "Reflections on Reflection: Van Fraassen on Belief," Synthese **98**, 297 (1994).

[26] B. C. van Fraassen, "Belief and the Problem of Ulysses and the Sirens," Phil. Stud. **77**, 7 (1995).

[27] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th anniversary edition, (Cambridge University Press, Cambridge, UK, 2010).

[28] M. Schlosshauer, *Decoherence and the Quantum-to-Classical Transition*, (Springer-Verlag, Berlin, 2007).