

Unscrambling Subjective and Epistemic Probabilities

Guido Bacciagaluppi

Abstract

There are two notions in the philosophy of probability that are often used interchangeably: that of subjective probabilities and that of epistemic probabilities. This paper suggests they should be kept apart. Specifically, it suggests that the distinction between subjective and objective probabilities refers to what probabilities *are*, while the distinction between epistemic and ontic probabilities refers to what probabilities are *about*. After arguing that there are *bona fide* examples of subjective ontic probabilities and of epistemic objective probabilities, I propose a systematic way of drawing these distinctions in order to take this into account. In doing so, I modify Lewis's notion of chances, and extend his Principal Principle in what I argue is a very natural way (which in fact makes chances fundamentally conditional). I conclude with some remarks on time symmetry, on the quantum state, and with some more general remarks about how this proposal fits into an overall Humean (but not quite neo-Humean) framework.

I believe I first met Itamar Pitowsky in what I think was the spring of 1993, when he came to spend a substantial sabbatical period at Wolfson College, Cambridge, while writing his paper on George Boole's 'Conditions of Possible Experience' (Pitowsky 1994). At that time, the Cambridge group was by far one of the largest research groups in philosophy of physics worldwide, and Wolfson had the largest share of the group. Among others, that included two further Israelis, my friends and fellow PhD-students Meir Hemmo, who is co-editing this volume, and Jossi Berkovitz, who had worked under Itamar's supervision on de Finetti's probabilistic subjectivism and its application to

quantum mechanics (published in English as Berkovitz 2012). We were all lunch regulars at Wolfson, and discussions on the philosophy of physics, of probability, and of mathematics were thus not limited to the setting of the Friday seminars in the HPS department (or the immediately preceding lunches at the Eraina Taverna).

My earliest introduction to de Finetti's subjectivism was in fact through Itamar and Jossi, and Itamar's seminal work on quantum probability (Pitowsky 1989, 1994) looms large within my formative years. The present paper relates most closely to Itamar's more recent work on Bayesian quantum probabilities (Pitowsky 2003, 2007). I shall not be agreeing with Itamar on everything, but that fits into the wider spirit of friendship, co-operation, integrity, and pursuit of knowledge that makes philosophers of physics an ideal model of an academic community, and Itamar an ideal model of a philosopher of physics. I dedicate this paper to his memory.

1 Subjective or epistemic probabilities?

This paper is about the notion of subjective probabilities and that of epistemic probabilities, and of how to keep them apart. Since terminology varies, I should briefly sketch which uses of the terms I intend. (A fuller explication will be given in Sections 2 and 3.)

I take the use of the term 'subjective probability' to be fairly uniform in the literature, and to be fairly well captured by the idea of *degrees of belief* (and as such 'assigned by agents' and presumably 'living in our heads'), as opposed to some appropriate notion of objective probabilities ('out there'), sometimes cashed out in terms of frequencies or propensities. I shall take the standard explication of the latter to be the one given by David Lewis with his notion of chance (Lewis 1985).

Along these lines, all probability assignments can be taken as subjective in the first place, but some such assignments may correspond to objective probabilities 'out there' in the world, determined by actual matters of fact. Probability assignments that do not correspond to objective probabilities shall be termed 'purely' or 'merely' subjective in the following. In this terminology,

the subjective probabilities of a subjectivist like de Finetti, for whom objective probabilities do not exist, are the prime example of purely subjective probabilities, even though usually they will be firmly if pragmatically rooted in matters of actual fact (as further discussed in Sections 4 and 10).

The term ‘epistemic probabilities’ has a wider variety of meanings, of which two are probably dominant, depending on whether one emphasises the fact that we lack knowledge of something or whether one emphasises the knowledge that we do have. The one I shall focus on is the former, that of *ignorance-interpretable* probabilities (i.e. of probabilities attaching to propositions expressing a matter of fact, but whose truth is unknown), which is perhaps most common in the philosophy of physics. The other refers to probability assignments whose *values* reflect our state of knowledge. The classical notion of probability based on the principle of indifference or an objective Bayesian view based on the principle of maximum entropy (to which we shall presently return) can thus also be termed epistemic.

This second notion of ‘epistemic probabilities’, which is perhaps most common in the philosophy of probability, belongs somewhere along the axis between subjective and objective probabilities. It is often contrasted with subjective probabilities when these are understood as merely ‘prudential’ or ‘pragmatic’ or ‘instrumental’. This distinction is challenged by Berkovitz (2019), who convincingly argues that de Finetti’s instrumental probabilities are also epistemic in this sense. Indeed, when talking about subjective probabilities in Sections 4, 6 and 10, I shall always implicitly take them to be so. This second sense of ‘epistemic probabilities’ is not the topic of this paper, and will be discussed no further.

Epistemic probabilities in the sense of ignorance-interpretable ones instead are liable to be assimilated to subjective probabilities, but, as I shall argue, they ought not to, because the distinction between probabilities that are epistemic in this sense and those that are not (we shall call them ‘ontic’) is conceptually orthogonal to the subjective-objective distinction.

Take for example Popper’s discussion in Vol. III of the *Postscript* (Popper 1982).¹ After treating extensively of his propensity interpretation in Vol. II, Popper starts Vol. III (on quantum mechanics) with a discussion of subjec-

¹Incidentally, reviewed by Itamar when it came out (Bub and Pitowsky 1985). Note that my examples of possible assimilation or confusion between epistemic and subjective

tive vs objective probabilities, complaining that physicists conflate the two (especially in quantum mechanics). Popper's contrast for the objective view of probabilities is what he throughout calls the subjective view, which for him is the view that probabilistic statements are about our (partial) *ignorance*. In the case of classical statistical mechanics, he complains that '[t]his interpretation leads to the absurd result that the molecules escape from our bottle because we do not know all about them' (Popper 1982, p. 109).

True, we care about whether our coffee goes cold, and the coffee does not care about our feelings. And, at least in the sense in which Popper took his own notion of propensities to be explanatory of observed frequencies (cf. Bub and Pitowski 1985, p. 541), it seems reasonable to say that while the behaviour of the coffee is a major factor in explaining why we come to have certain subjective expectations, it is only objective probabilities that can explain the behaviour of the coffee. But there is also a sense in which the probabilities of (classical) statistical mechanics are *clearly* epistemic: we do believe the coffee or the gas in the bottle to be in a definite microstate, but we do not know it. Also the statistical mechanical description of the gas presupposes there to be a particular matter of fact about the state of the molecules, which is however left out of the description. If we read Popper's complaint as being that epistemic probabilities *must* be subjective, or rather purely subjective (i.e. unrelated to objective probabilities), then it seems that *no* explanation of the kind Popper envisages is possible at all in statistical mechanics.

Or take Jaynes, e.g. in 'Some Random Observations' (1985). For Jaynes all probabilities are fundamentally subjective in the sense that 'a *probability* is something that we *assign* in order to represent a state of knowledge', as opposed to an actual frequency, which is a 'factual property of the real world' (p. 120). While for Jaynes the probabilities of classical statistical mechanics are thus of course epistemic, unlike for Popper they are meant to be (at least partially?) *objective*, in the sense that there are (rational?) principles, first and foremost the principle of maximum entropy, that provide a way for

probabilities (in the senses just sketched) may well prove to be straw-mannish under further scrutiny, but I just wish to make a *prima facie* case that these two notions can indeed be confused. The rest of the paper will make the extended case for why we should distinguish them and how.

choosing our probability assignments based on those facts we do know.²

Part of Jaynes's motivation came from the complaint that the mathematics of quantum theory 'describes in part physical law, in part the process of human inference, all scrambled together in such a way that nobody has seen how to separate them' (p. 123). He continues:

Many years ago I became convinced that this unscrambling would require that probability theory itself be reformulated so that it recognizes explicitly the role of human information and thus restores the distinction between reality and our knowledge of reality, that has been lost in present quantum theory.³

Jaynes' position, however, faces a different problem precisely in the case of quantum mechanics. Whether or not probabilities are purely subjective, for Jaynes they are always epistemic: '[p]robabilities in present quantum theory express the incompleteness of human knowledge just as truly as did those in classical statistical mechanics' (1985, p. 122). And this ominously leads him to question the possibility of fundamental indeterminism.

²It is unclear to what extent this 'objective Bayesian' strategy indeed yields objective probabilities. Consider for instance situations in which information about the physical situation is in principle readily available but is glibly ignored. In this case, even if we use principles such as maximum entropy, one might complain that there is little or no connection between the physical situation and the probabilities we assign. On the other hand, there may be situations in which the physical situation itself appears to systematically lead to situations of ignorance, e.g. because of something like mixing properties of the dynamics. In this case, there may arguably be a systematic link between certain objective physical situations and certain ways of assigning probabilities, which may go some way towards justifying the claim that these probabilities are objective. (Cf. discussions about whether high-level probabilities are objective, e.g. Glynn (2009), Lyon (2011), Emery (2013) and references therein.) If so, note that this will be the case whether or not there is *anyone* to assign those probabilities. Many thanks to Jossi Berkovitz for discussion of related points.

³For an example of what Jaynes has in mind, see Heisenberg's 'The Copenhagen interpretation of quantum theory' (Heisenberg 1958, Chap. 3). In Section 9 below, I shall allude to what I think are the origins of these aspects of Heisenberg's views, but I agree with Jaynes that (at least in the form they are usually presented) they are hard to make sense of. Note also that 'the Copenhagen interpretation' as publicised there appears to have been largely a (very successful) public relations exercise by Heisenberg (cf. Howard 2004). In reality, physicists like Bohr, Heisenberg, Born and Pauli held related but distinct views about how quantum mechanics should be understood.

Indeed, as Jaynes sees it, in the case of quantum mechanics human knowledge is limited by dogmatic adherence to indeterminism in the Copenhagen interpretation. But in fact, even though the ‘Copenhagen interpretation’ may have been dogmatic, indeterminism was the least worrying of its dogmas: a wide spectrum of non-dogmatic approaches to the foundations of quantum mechanics today embrace indeterminism in some form or other. Thus, if (as Jaynes seems to think) subjective probabilities *must* be epistemic, then they seem to be inapplicable to many if not most fundamental approaches to quantum mechanics.⁴

These problems with both Popper’s and Jaynes’ positions seem to me to stem from a scrambling together of subjective and epistemic probabilities. I shall take it as uncontroversial that we often form degrees of belief about facts we do not know (so that such probabilities are both subjective and epistemic), but this in itself does not show that all epistemic probabilities should be subjective, or that all subjective probabilities should be epistemic. Indeed, being a degree of belief is conceptually distinct from being a probability about something we ignore: the characterisation of a probability as subjective is in terms of what that probability *is*, that of a probability as epistemic is in terms of what that probability is *about*.

Admittedly, even if we can conceptually distinguish between subjective and epistemic probabilities, it could be that the two categories are coextensive. And if all epistemic probabilities are merely subjective, or all subjective probabilities are also epistemic, then Popper’s or Jaynes’s worries become legitimate again. As a matter of fact, typical cases of epistemic probabilities are arguably subjective (‘Is Beethoven’s birthday the 16th of December?’⁵), and, say, probabilities in plausibly indeterministic scenarios are typically

⁴Spontaneous collapse approaches are invariably indeterministic (unless one counts among them also non-linear modifications of the Schrödinger equation, which however have not proved successful), as is Nelson’s stochastic mechanics. Everett’s theory vindicates indeterminism at the emergent level of worlds. And even many pilot-wave theories are stochastic, namely those modelled on Bell’s (1984) theory based on fermion number density. Thus, among the main fundamental approaches to quantum mechanics, only theories modelled on de Broglie and Bohm’s original pilot-wave theory provide fully deterministic underpinnings for it. (For a handy reference to philosophical issues in quantum mechanics, see Myrvold (2018b) and references therein. For more on Nelson see footnote 40.)

⁵16 December is the traditional date for Beethoven’s birthday, but we know from documentary evidence only that he was baptised on 17 December 1770 – Beethoven was *probably* born on 16 December.

thought to be objective (‘Will this atom decay within the next minute?’). And, indeed, while there is an entrenched terminology for probabilities that are not purely subjective (namely ‘objective probabilities’), non-epistemic probabilities do not even seem to have a separate standard name (we shall call them ‘ontic probabilities’). It is indeed tempting to think that all epistemic probabilities are subjective and all ontic probabilities are objective.

I urge the reader to resist this temptation. I believe that we can make good use of the logical space that lies between these distinctions. As the Popper and Jaynes quotations above are meant to suggest, objectivism about probability will be self-limiting if it takes itself to apply only to ontic probabilities, as will subjectivism about probability if it takes itself to apply only to epistemic probabilities. The landscape of the philosophy of probability will be both richer and more easily surveyable if we unscramble subjective and epistemic probabilities in a way that allows also for (purely) subjective ontic probabilities as well as for epistemic objective probabilities. While this is surely not the first paper to make some of these claims (indeed, Jaynes is obviously proposing that epistemic probabilities can be objective⁶), I believe the way it approaches these questions will be original.

In what follows, I shall first recall the standard discussion of subjective and objective probabilities given by David Lewis (Section 2), then try and spell out what one typically means by the distinction between epistemic and ontic probabilities (Section 3). This will provide enough background to argue that there are *bona fide* cases of probabilities that are both purely subjective and ontic (Section 4), and of probabilities that are both epistemic and objective (Section 5). This will lead to an extension of Lewis’s definition of objective probabilities and of his Principal Principle, in which chances become fundamentally conditional, thus redrawing the subjective-objective distinction (Section 6), and to a more systematic discussion also of the epistemic-ontic distinction, clarifying some further ambiguity (Section 7). I conclude the paper with some additional remarks on time symmetry and on the quantum state (Sections 8 and 9),⁷ and with more general remarks on how the pro-

⁶With the qualifications mentioned in footnote 2. Thanks to both Jos Uffink and Ronnie Hermens for making me include more about Jaynes.

⁷These (especially the latter) can be skipped by readers who are not particularly familiar with debates on the metaphysics of time or, respectively, on the nature of the quantum state and of collapse in quantum mechanics.

posed framework for conceptualising probabilities fits into an overall Humean framework – at the same time distancing myself from some typical neo-Humean views (Section 10).

2 Subjective vs objective probabilities

Recall the standard distinction between *subjective* and *objective* probabilities: subjective probabilities are coherent degrees of belief, generally taken to live in our head, while objective probabilities are meant to live out there in the world, properties of external objects or reduced to such properties. I shall follow (half of) the literature in taking subjective probabilities as well-understood, and objective probabilities (if they exist) as being the more mysterious concept. While the literature on subjective probabilities is considerable and varied, I shall follow what is nowadays often considered the standard analysis of how subjective and objective probabilities relate – and which in fact provides an implicit definition of objective probabilities given an understanding of subjective probabilities – David Lewis’s ‘A Subjectivist’s Guide to Objective Chance’ (Lewis 1980).⁸

Importantly for our later discussion, Lewis takes objective probabilities (also called ‘chances’) to be functions of propositions and times. Assuming for the sake of example that there are chances for rolls of dice, before I roll a pair of dice we have the intuition that the chance of rolling 11 is $1/18$. Once I have rolled the first die, the chance of 11 either goes down to 0 and stays there (if I have rolled one of the numbers 1 to 4), or it goes up to $1/6$ (if I have rolled 5 or 6). Then when I roll the second die the chance of 11 goes definitively either down to 0 or up to 1.

The idea behind Lewis’s analysis is that objective probabilities, if they exist, are the kind of thing that – in a sense to be made precise – would rationally constrain our subjective probabilities (also called ‘credences’⁹) *completely*.

⁸While in a Cambridge HPS mood, I cannot refrain from thanking Jeremy Butterfield for (among so many other things!) introducing me to this wonderful paper and all issues around objective probabilities.

⁹Again, terminology varies somewhat, but I shall take ‘credences’ to be synonymous with ‘subjective probabilities’.

Credences of course are standardly taken to be rationally constrained both synchronically by the probability calculus and diachronically by Bayesian conditionalisation.¹⁰ Few other requirements are widely recognised as able to constrain their values further (maybe only the requirement that credences should be non-dogmatic, i. e. not take the values 0 or 1, is one commanding near-universal support – but see the remarks below in Section 4). Thus, the criterion that Lewis proposes for objective probabilities is very stringent.

What Lewis has in mind is the following. The chance at t of A is x iff there is a proposition $\text{ch}_t(A) = x$ such that (subject to a few provisos to be spelled out) conditionalising on $\text{ch}_t(A) = x$ will always yield the *unique* value x for our credence,

$$\text{cr}_t(A \mid \text{ch}_t(A) = x) = x, \quad (1)$$

whatever the form of our credence function at the time t . Of course, if we further believe at t that the chance at t of A is in fact x , i. e. if $\text{cr}_t(\text{ch}_t(A) = x) = 1$, then our credence at t in A is also x (by the probability calculus), or if we learn at t that $\text{ch}_t(A) = x$, our credence in A becomes x (by Bayesian conditionalisation).

The stringency of Lewis’s requirements on what is to count as rationally compelling is such that there are few candidates in the literature for satisfying them, none of which are uncontroversial. And because Lewis himself believes in ‘Humean supervenience’, for him chances have to supervene on ‘particular matters of fact’, making it especially hard to find plausible explicit candidates for objective chances. This problem is quaintly known in the literature as the ‘big bad bug’.¹¹

¹⁰Bayesian conditionalisation is often taken as the unique rational way in which one can change one’s credences. But there may be other considerations competing with it, in particular affecting background theoretical knowledge. This will generally be implicit (but not needed) in much of this paper. Thanks to Jossi Berkovitz for pointing out to me that de Finetti himself was not committed to conditionalisation as the only way for updating probabilities, believing instead that agents may have reasons for otherwise changing their minds (cf. Berkovitz 2012, p. 16).

¹¹Two proposals of note are the one by Deutsch and by Wallace in the context of Everettian quantum mechanics, who derive the quantum mechanical Born rule from what they argue are rationality constraints on decision making in the context of Everett (see e.g. Wallace (2010), or Brown and Ben Porath (this volume)), and the notion of ‘Humean objective chances’ along the lines of the neo-Humean ‘best systems’ approach (Loewer 2004, Frigg and Hoefer 2010). Note that I am explicitly distinguishing between Lewisian

The provisos are, first, that our credence function at t should not rule out that $\text{ch}_t(A) = x$, i. e. we should have $\text{cr}_t(\text{ch}_t(A) = x) \neq 0$, otherwise the credence in (1) will be ill-defined; second, and more subtly, that our credence function at t should not incorporate information that is ‘inadmissible’ at t . To make this intuitive, think of the chance at t of A as fixing the objectively fair odds for betting on A . Clearly, if in advance of making the bet either you or the bookie had some more direct information about whether or not A will come to pass (e. g. one of you is clairvoyant), that would be cheating. In order for the notion of ‘objectively fair odds’ at time t to make sense, we need to have an appropriate notion of what is admissible in shaping our credences at t . Not everything can be admissible, for then we would allow A , or not- A , and the only propositions about chance that could satisfy (1) would correspondingly be propositions implying $\text{ch}_t(A) = 1$, or propositions implying $\text{ch}_t(A) = 0$, and all chances would be trivial. On the other hand, some propositions must be admissible, since $\text{ch}_t(A) = x$ clearly is. We can write our credence function $\text{cr}_t(A)$ as $\text{cr}(A|E_t \& T)$, where $\text{cr}(A)$ is a ‘reasonable initial credence function’, T represents any background assumptions that we make (which might perhaps better go in a subscript), and E_t is our epistemic context at time t (by which I mean all propositions about matters of fact known at time t). If we leave fixed the background assumptions T , then $\text{cr}_t(A)$ just evolves by Bayesian conditionalisation upon the new facts that we learn between any two times t and t' .¹²

chances as implicitly defined by the PP, and Humean objective chances. The latter are an attempt to find something that will satisfy the definition of Lewisian chances, but there is no consensus as to whether it is a successful attempt. (See Section 10 below for some further comments on neo-Humeanism.) The Deutsch-Wallace approach is often claimed to provide the only known example of Lewisian chances, but again this claim is controversial, as it appears that some of the constraints that are used cannot be thought of as rationality assumptions but merely as naturalness assumptions. On this issue see Brown and Ben Porath (this volume) as well as the further discussion in Saunders *et al.* (2010). The difficulty in finding anything that satisfies the definition of Lewisian chances is (of course) an argument in favour of subjectivism.

¹²I shall assume throughout that *degrees of belief* attach to propositions about matters of fact (in the actual world), which incidentally need not be restricted to Lewis’s ‘particular’ matters of fact (e.g. I shall count as matters of fact also holistic facts about quantum states – if such facts there be). I shall call these *material* or *empirical* propositions.

Of course, some of the background assumptions that we make in choosing our degrees of belief may not just refer to matters of fact, but will be *theoretical* propositions (say, referring to possible worlds). In that case, I would prefer not to talk about degrees of belief attaching to such theoretical propositions, but about *degrees of acceptance*.

Substituting into (1), we have

$$\text{cr}\left(A \mid (\text{ch}_t(A) = x) \ \& \ E_t \ \& \ T\right) = x . \quad (2)$$

And the requirement that (1) should hold whatever the form of our credence function at t (subject to the provisos) now translates into the requirement that (2) should hold for any ‘reasonable initial credence function’ (which we shall keep fixed) and for any admissible propositions E_t and T , which for Lewis are, indeed, ones that are compatible with the chance of A being x and in addition do not provide information about the occurrence of A over and above what is provided by knowing the chance of A (if chance is to guide credence, we must not allow information that trumps the chances). This is

The matter is further complicated by the fact that even full acceptance of a theory need not imply belief in all propositions about matters of fact following from the theory. For instance, an empiricist may believe only a subset of the propositions about actual matters of fact following from a theory they accept (say, only propositions about observable events), thus drawing the empirical-theoretical distinction differently and making a distinction between material propositions and empirical propositions, which for the purposes of this paper I shall ignore. (For my own very real empiricist leanings, see Bacciagaluppi (2019).)

For simplicity I shall also ignore the distinction between material propositions and actual propositions that one makes in a many-worlds theory, where the material world consists of several ‘possible’ branches, even though I shall occasionally refer to Everett for illustration of some point.

One might reflect the distinction between material (or empirical) and theoretical propositions by not conditionalising on the latter, but letting them parameterise our credence functions. Another way out of these complications may be to assign subjective probabilities to all propositions alike, but interpret the probabilities differently depending on what propositions they attach to: in general they will be degrees of acceptance, but in some cases (all material propositions, or a more restricted class of empirical propositions), one may think of them as straightforward degrees of belief. Such an additional layer of interpretation of subjective probabilities will, however, not affect the role they play in guiding our actions (degrees of acceptance will generally be degrees of ‘as-if-belief’). In any case, the stroke notation makes clear that assuming a theory – even hypothetically – is meant to affect our credences as if we were conditionalising on a proposition stating the theory (I believe this is needed to derive some of Lewis’s results).

Note that for Lewis, T will typically include theoretical propositions about chances themselves in the form of history-to-chance conditionals, and because in Lewis’s doctrine of Humean supervenience chances in fact supervene on matters of fact in the actual world, such background assumptions may well be propositions that degrees of belief can attach to (at least depending on how one interprets the conditionals), even if one wishes to restrict the latter (as I have just sketched) to propositions about matters of fact in the actual world.

We shall return to the material-theoretical distinction in Section 7.

Lewis’s ‘Principal Principle’ (PP) relating credence and chance. It can be taken, is often taken, and I shall take it as providing the implicit definition of objective probabilities: iff there is a proposition X such that, for all choices of reasonable initial credence function and admissible propositions,

$$\text{cr}(A \mid X \& E_t \& T) = x , \quad (3)$$

then there is such a thing as the objective chance at t of A and its value is x . The weakest such proposition X is then $\text{ch}_t(A) = x$.¹³

More precisely, Lewis restricts admissible propositions about matters of fact to historical propositions up to time t (at least ‘as a rule’: no clairvoyance, time travel and such, and we shall follow suit for now – although he appears to leave it open that one may rethink the analysis if one should wish to consider such less-than-standard situations (Lewis 1980, p. 274)). Admissible propositions about matters of fact can thus be taken to be what can be reasonably thought to be knowable in principle at t , and Lewis allows them to include the complete history H_t of the world up to that time. He also allows hypothetical propositions about the theory of chance, i. e. about how chances may depend on historical facts.

Even more propositions might be admissible at time t , but this partial characterisation of admissibility already makes the PP a useful and precise enough characterisation of objective chances. Indeed, taking the conjunction $H_t \& T_C$ of the complete history of the world up to t and a complete and correct theory T_C of how chances depend on history, the following will be an instance of the PP:

$$\text{cr}(A \mid (\text{ch}_t(A) = x) \& H_t \& T_C) = x . \quad (4)$$

Since T_C is a complete and correct theory of chance, $H_t \& T_C$ already implies that $\text{ch}_t(A) = x$, and we can rewrite (4) as

$$\text{ch}_t(A) = \text{cr}(A \mid H_t \& T_C) . \quad (5)$$

This instance of the PP is used by Lewis to derive various properties of chances using known properties of credences, e. g. that they satisfy the probability calculus, and that they evolve by conditionalisation on intervening

¹³Note that X is a proposition about (actual) matters of fact. Thus, chances in the sense of Lewis indeed satisfy the intuition that objective probabilities are properties of objects or more generally determined by actual facts. Thanks to Jossi Berkovitz for discussion.

events. In particular, according to Lewis, chances of past events are all 0 or 1. Taking $\text{ch}(A)$ as the chance function at some arbitrary time $t = 0$, we can thus for all later times write $\text{ch}_t(A) = \text{ch}(A|H_t)$.

3 Epistemic vs ontic probabilities

A systematic treatment of the epistemic-ontic distinction comparable to Lewis's treatment of subjective and objective probabilities is not available. But we can get a reasonably good idea of such a distinction, which will be enough for the time being, by looking at common ways of using the term 'epistemic probabilities'. (We shall revisit this distinction in Section 7.)

Epistemic probabilities are probabilities that attach to propositions of which we are ignorant. An obvious example is: I flip a coin and do not look at the result yet – what is the probability of Heads? In this situation, there is a matter of fact about whether the coin has landed Heads or Tails, and if I knew it, I would assign to Heads probability 0 or 1 (depending on what that matter of fact was).

Another common example is modelling coin flipping as a deterministic Newtonian process: even before the coin has landed, the details of the flip (initial height, velocity in the vertical direction, and angular momentum¹⁴) determine what the outcome will be, but we generally do not know these details. In this case, we say that the probabilities of Heads and Tails are epistemic, even if the event of the coin landing Heads or Tails is in the future, because we can understand such events as labels for sets of unknown initial conditions.

If we are looking for examples of ontic probabilities, we thus presumably need to look at cases of indeterminism. The most plausible example is that of quantum probabilities, say, the probability for a given atom to decay within the next minute. (Of course, whether this is a genuine case of indeterminism will depend on one's views about quantum mechanics, in particular about the existence of 'hidden variables'. For the sake of the example we shall take quantum mechanics to be genuinely indeterministic.)

¹⁴For the classic Newtonian analysis of coin flipping see Diaconis, Holmes and Montgomery (2007).

Still, if all actual matters of fact past, present and future are taken to be knowable in principle, then all probabilities become purely epistemic, and (at least for now) it appears we cannot draw an interesting distinction.

One way of ensuring that the future is unknowable, of course, is to postulate that the future is *open*, i. e. that at the present moment there is no (single) future yet. This is a very thick metaphysical reading of indeterminism, however, which I wish to avoid for the purpose of drawing the epistemic-ontic distinction.¹⁵ Indeterminism can be defined in a non-metaphysical way in terms of the allowable models of a theory (or nomologically possible worlds, if worlds are also understood in a metaphysically thin way): a theory is (future-)deterministic iff any two models agreeing up to (or at) time t agree also for all future times (and with suitable modifications for relativistic theories).

For the purpose of this first pass at drawing the epistemic-ontic distinction we can then simply stipulate that the notion of ignorance only pertains to facts in the actual world that are knowable in principle, and that at time t only facts up to t are knowable in principle (which is in fact the stipulation we already made in the context of Lewis's analysis of chances – we may want again to make a mental note to adapt the notion of 'knowable in principle' if we allow for clairvoyance, time travel and such like). If one so wishes, one can of course talk about 'ignorance' of which possible world is the actual world, or of which way events in our world will be actualised. But if probabilities in cases of indeterminism are to count as 'ontic', then it simply is the case that in practice we make a more restrictive use of the term 'ignorance' when we draw the epistemic-ontic distinction.

The general intuition is now that epistemic probabilities lie in the *gap* between probability assignments given what is knowable in principle (call it the *ontic context*), and probability assignments given what is actually known (the

¹⁵Or for any other purpose – as a matter of fact, I shall assume throughout a block universe picture. Specifically, in the case of indeterminism one should think of *each possible world* as a block universe. The actual world may be more real than the merely possible ones, or all possible worlds may be equally real (whether they are or not is inessential to the following, although a number of remarks below may make it clear that I am no modal realist, nor in fact believe in objective modality). But *within* each world, all events will be equally real. (Note that at a fundamental level, the Everett theory is also deterministic and Everettians emphasise that the Everettian multiverse is indeed a block universe.)

epistemic context). By the same token, ontic probabilities are any probabilities that remain non-trivial if we assume that we know everything that is knowable in principle.

In general probabilities will be a *mix* of epistemic and ontic probabilities. For instance, we ask for the probability that any atom in a given radioactive probe may decay in the next minute, but the composition of the probe is unknown: the probability of decay will be given by the ontic probabilities corresponding to each possible radioactive component, *weighted* by the epistemic probabilities corresponding to each component.

The EPR setup in quantum mechanics (within a standard collapse picture) provides another good example: before Bob performs a spin measurement, assuming the global state is the singlet state, Alice has ontic probability 0.5 of getting spin up in any direction. After Bob performs his measurement, Alice has ontic probabilities $\cos^2(\alpha/2)$ and $\sin^2(\alpha/2)$ of getting up or down, where α is the angle between the directions measured on the two sides. However, she is ignorant of Bob's outcome, and her epistemic probabilities for his having obtained up or down are both 0.5. Thus, given her epistemic context, her (now mixed) probabilities for her own outcomes are unchanged.¹⁶

As a final example, assume we have some ontic probability that is also an objective chance in the sense of Lewis. According to the PP in the form (5), this ontic probability will be given by

$$\text{ch}_t(A) = \text{cr}(A|H_t \& T_C) . \quad (6)$$

Now compare it with $\text{cr}(A|E_t \& T_C)$, our credence in A at t given the same theoretical assumptions but our actual epistemic context E_t . This now is a mixed (partially epistemic) probability. Indeed, let H_t^i be the various epistemic possibilities (compatible with our epistemic context E_t) for the actual history up to t (i. e. for the actual ontic context). We then have

$$\text{cr}(A|E_t \& T_C) = \sum_i \text{cr}(A|H_t^i \& E_t \& T_C) \text{cr}(H_t^i|E_t \& T_C) , \quad (7)$$

¹⁶This is of course the quantum mechanical no-signalling theorem, as seen in the standard collapse picture. (It is instructive to think through what the theorem means in pilot-wave theory or in Everett. Thanks to Harvey Brown for enjoyable discussion of this point.) And of course, mixed quantum states are the classic example of the distinction in quantum mechanics between probabilities that are ignorance-interpretable ('proper mixtures') and probabilities that are not ('improper mixtures').

which simplifies to

$$\text{cr}(A|E_t \& T_C) = \sum_i \text{cr}(A|H_t^i \& T_C) \text{cr}(H_t^i|E_t \& T_C) , \quad (8)$$

because the H_t^i are fuller specifications of the past history at t than the one provided by E_t . In (8), now, the credences $\text{cr}(H_t^i|E_t \& T_C)$ are purely epistemic probabilities, and the $\text{cr}(A|H_t^i \& T_C)$ are purely ontic probabilities corresponding to each of the epistemically possible ontic contexts. The various H_t^i could e. g. fix the composition of the radioactive probe in our informal example above.¹⁷

Judging by the examples so far, the subjective-objective distinction and the epistemic-ontic distinction could well be co-extensive. Indeed, both subjective and epistemic probabilities seem to have to do with us, while both objective and ontic probabilities seem to have to do with the world. Epistemic probabilities depend on epistemic contexts, as do subjective ones; and chances depend on past history, which we have also taken as the context in which ontic probabilities are defined. But appearances may deceive.

In order to argue that the subjective-objective distinction and the epistemic-ontic distinction are to be drawn truly independently, I need to convince you that there are *bona fide* examples of probabilities that are both subjective (in fact purely subjective) and ontic, and of probabilities that are both epistemic and objective. These will be the respective objectives of the next two sections.

4 Subjective ontic probabilities

Let us start with the case of subjective ontic probabilities. I have discussed this case before, in a paper (Bacciagaluppi 2014) in which I comment on

¹⁷In this example, it is of course possible to think of mixed epistemic and ontic probabilities as arising not only through ignorance of matters of fact in the ontic context H_t , but also through ignorance of the correct theory of chance T_C . We might indeed know the exact composition of the radioactive probe, but take ourselves to be ignorant of the correct decay laws for the various isotopes, with different credences about which laws are correct. I shall neglect this possibility in the following, because as noted already in footnote 12, I prefer not to think of ignorance when talking about theoretical assumptions. In any case, I shall not need this further possible case for the arguments below.

Chris Fuchs and co-workers' radical subjectivist approach to quantum mechanics, known as qBism (see e.g. Fuchs 2002, 2014),¹⁸ and in which I spell out how I believe radical subjectivism *à la* de Finetti (1970) can be applied to quantum mechanics not only along qBist lines but also within any of the traditional foundational approaches (Bohm, GRW and Everett). The idea that subjectivism is applicable to quantum mechanics was in fact prompted by Itamar, and worked out by Jossi Berkovitz in the years 1989–1991 with specific reference to de Finetti's idea that probabilities are only definable for verifiable events.

For the purposes of the present paper, I need to make the case that *purely* subjective probabilities can be ontic, thus radical subjectivism about quantum probabilities (in a context in which quantum mechanics is taken to be irreducibly indeterministic) provides the ideal example, and in the following I shall rely heavily on my own previous discussion. (Alternatively, I could arguably have relied on Berkovitz (2012).) Note that what I need to convince you of is not that such a position is correct, but only that it is coherent, i. e. that there is an interesting sense in which one can combine the notion of purely subjective probabilities with that of ontic probabilities.

Itamar himself (2003, 2007), Bub (2007, 2010) and Fuchs and co-workers have all defended approaches to quantum mechanics in which probabilities are subjective in the sense of being degrees of belief, and some modifications or additions to Bayesian rationality are introduced in order to deal with situations, respectively, in which one deals with 'quantum bets' (incompatible bets possibly with some outcomes in common), in which one assumes no-cloning as a fundamental principle, or in which gathering evidence means irreducibly intervening into the world. I take it that all of these authors (as well as de Finetti himself) agree on the assumption that in general there are no present or past facts that will determine the outcomes of quantum bets or any quantum measurements, and thus that the probabilities in their approaches are ontic in the sense used in the previous section.

It is less clear that the probabilities involved are purely subjective. Indeed, both Pitowsky and Bub see subjective probability assignments as severely constrained by the structure of quantum bets or quantum observables, in

¹⁸I am deeply grateful to Chris Fuchs for many lovely discussions about qBism and subjectivism over a number of years.

particular through Gleason’s theorem and (relatedly) through the mathematical structure of collapse. These might be seen as rationality constraints, so that the probabilities in these approaches might in fact be seen as at least partly objective.¹⁹ As Itamar puts it (Pitowsky 2003, p. 408):

For a given single physical system, Gleason’s theorem dictates that all agents share a common prior probability distribution or, in the worst case, they start using the same probability distribution after a single (maximal) measurement.

Indeed, since pure quantum states always assign probability 1 to results of some particular measurement, the constraints may be so strong that even some dogmatic credences are forced upon us.

Fuchs instead clearly wishes to think of his probabilities as radically subjective (which I want to exploit). Precisely because of the strong constraints imposed by collapse, he feels compelled to not only take probability assignments as subjective, but also quantum states themselves and even Hamiltonians (Fuchs 2002, Sect. 7). To a certain extent, however, I think this is throwing out the baby without getting rid of the bath water. Indeed, I believe the (empirically motivated but normative) additions to Bayesian rationality to which Fuchs is committed imply already by themselves very strong constraints on qBist probability assignments (Bacciagaluppi 2014, Sect. 2.2).

¹⁹Whether or not they are will make little difference to my argument (although if they should be purely subjective I would have further concrete examples of subjective ontic probabilities in the sense I need). In Bub’s case, he explicitly states that the no-cloning principle is an assumption and could turn out to be false. That means that the constraints hold only given some specific theoretical assumptions, and are not rationally compelling in the strong sense used by Lewis. In Pitowsky’s case, the constraints seem stronger, in the sense that the ‘quantum bets’ are not *per se* theoretical assumptions about the world, but are just betting situations whose logical structure is non-classical. Whether or not there are rational constraints on our probability assignments will be a well-posed question only *after* you and the bookie have agreed on the logical structure of the bet. On the other hand, whenever we want to evaluate credences in a particular situation, we need to make a theoretical judgement as to what the logic of that situation is. That means, I believe, that whether we can have rationally compelling reasons to set our credences in any particular situation will still depend on theoretical assumptions, and fall short of Lewis’s very stringent criterion. See Bacciagaluppi (2016) for further discussion of identifying events in the context of quantum probability, and Sect. 1.2 of Pitowsky (2003) for Itamar’s own view of this issue. Cf. also Brown and Ben Porath (this volume).

On the other hand, as I shall presently explain, I also believe that radical subjectivism can be maintained even within a context in which quantum states and Hamiltonians are taken as (theoretical) elements in our ontology. While in some ways less radical than Fuchs’s own qBism, this is the position that I propose in my paper, and which I take to prove by example that a coherent case can be made for (purely) subjective ontic probabilities. It can be summarised as follows.

Subjectivism about probabilities (de Finetti 1970) holds that there are *no* objective probabilities, only subjective ones. Subjective probability assignments are strategies we use for navigating the world. Some of these strategies may turn out to be more useful than others, but that makes them no less subjective. As an example, take again (classical) coin flipping: we use subjective probabilities to guide our betting behaviour, and these probabilities get updated as we go along (i.e. based on past performance). But it is only a contingent matter that we should be successful at all in predicting frequencies of outcomes. Indeed, *any* given sequence of outcomes (with any given frequency of Heads and Tails) will result from appropriate initial conditions.

Under certain assumptions the judgements of different agents will converge with subjective probability 1 when conditionalised on sufficient common evidence, thus establishing an intersubjective connection between degrees of belief and frequencies. This is de Finetti’s theorem (see e.g. Fuchs 2002, Sect. 9). One might be tempted to take it as a sign that subjective probabilities are tracking something ‘out there’.²⁰ This, however, does not follow – because the assumption of ‘exchangeability’ under which the theorem can be proved is itself a subjective feature of our probability assignments. The only objective aspect relevant to the theorem is the common evidence, which is again just *past* performance.

For the subjectivist, there is no sense in which our probability judgements are right or wrong. Or as de Finetti (1970, Preface, quoted in Fuchs 2014) famously expressed it: ‘PROBABILITIES DO NOT EXIST’. While this sounds radical, it is indeed perfectly unremarkable as long as we assume that states of affairs are always determinate and evolve deterministically. In that case, Lewis himself will say that objective chances are all 0 or 1, so that

²⁰Such a reading is e.g. perhaps present in Greaves and Myrvold (2010), or even in Pitowsky (2003).

all non-trivial probabilities are indeed subjective and (non-trivial) chances do not exist.

What we need to consider is the case of indeterminism. Surely the situation is completely different in quantum mechanics? As Lewis puts it in the introduction to the ‘Subjectivist’s Guide’, ‘the practice and analysis of science require both [subjective and objective probabilities]’. He accordingly demands that also subjectivists should be able to make sense of the proposition that ‘any tritium atom that now exists has a certain chance of decaying within a year’ (Lewis 1980, p. 263). Thus, we are told, de Finetti’s eliminative strategy towards chance cannot be applied to quantum mechanics.

But, actually, why not? In the deterministic case, de Finetti’s analysis shows us how we can have a useful and even intersubjectively robust notion of probability that relates to actual matters of fact but is conceptually quite autonomous from them. And after all, it is especially when we take determinism to fail and the world around us to behave randomly that we may most need (probabilistic) strategies to navigate the world.

What goes unappreciated in the sweeping intuition that subjectivism must fail in the case of indeterminism is precisely the role and status of the *strategies* we apply in choosing our probability assignments. Subjectivists may well use certain strategies or shortcuts to decide what subjective probability assignments to make. These are just what we called *theoretical assumptions* or *models* in Section 2. They are not Lewisian rationality constraints on probability assumptions, but systematic ways in which we subjectively choose our credences. Data may well ‘confirm’ such models, but again that just means that our subjective probability assignments have performed well in the past.

The form theoretical models may typically take is that they will systematically connect certain non-probabilistic features of a given situation with certain probability assignments. For instance, we may choose to assign certain subjective probabilities to rolls of dice given their internal weight distribution. For the subjectivist, these probabilities are and remain purely subjective. Indeed, there is no necessary connection between the weight distribution in a die and the outcomes of our rolls, not just in the Humean sense that these are independent matters of fact, but even in the sense that, given a Newtonian model of dice rolling, for any possible sequence of results there are initial conditions that will produce it. Thus, there is no Lewisian com-

elling rational justification for letting our credences depend on the weight distribution in the die, only pragmatic criteria such as simplicity, again past performance, etc.

From this perspective, *quantum mechanics* is just such a theoretical model for choosing our subjective probabilities, and indeed a very successful one. One can take the analogy with probabilistic models of dice rolling quite far: unlike Fuchs, who *identifies* quantum states with catalogues of subjective probabilities, I believe a subjectivist may take the quantum state itself as a non-probabilistic, ontic feature of the world and use it (like weight distribution) as the basis for choosing our subjective probabilities.²¹ There is no need for a qBist like Fuchs to reject the putative onticity of quantum states,²² precisely because that by itself does not make our probability assignments objective (forgive me, Itamar!).

As further described in Bacciagaluppi (2014), this general perspective can be applied to all the main approaches to quantum mechanics that take the quantum state as ontic, so that one can be a de Finettian subjectivist about probabilities in the context of Bohm, GRW or Everett. In the case of Bohm, this is unremarkable of course, because the Bohm theory is deterministic and probabilities are epistemic in the sense used above. In the case of GRW, it is an unfamiliar position. The default intuition in GRW is presumably to take probabilities as some kind of propensities, but subjectivism may prove more attractive.²³ Finally, the Deutsch-Wallace decision-theoretic framework for the Everett theory already analyses probabilities as subjective. Deutsch and Wallace themselves (see e.g. Wallace 2010) argue that the weights in the universal wavefunction in fact *satisfy* the Principal Principle (and are in fact

²¹The quantum state is thus taken to be a matter of fact, but presumably not a ‘particular matter of fact’ in the sense of Lewis, because of the holistic aspects of the quantum state. (Cf. footnote 12 above.) Note also that historically the notion of quantum state *preceded* its ‘statistical interpretation’, famously introduced by Max Born (for discussion see Bacciagaluppi 2008).

²²For other, even more forceful arguments to this conclusion, see Brown (forthcoming) and Myrvold (this volume).

²³Frigg and Hoefer (2007) have argued that the propensity concept is inadequate to the setting of GRW. Their own preference is to use ‘Humean objective chances’ instead (cf. Frigg and Hoefer 2010). Note that if the alternative subjectivist strategy includes taking the quantum state as a non-probabilistic physical object, there are severe constraints on doing this for the case of *relativistic* spontaneous collapse theories (see also below, Sections 6 and 9).

the only concrete example of something that qualifies as objective chances in Lewis’s sense!). Many authors, however, dispute that the Deutsch-Wallace theorem rationally compels one to adopt the Born rule. I concur, but see that as no reason to reject the decision-theoretic framework itself. One can simply use the Deutsch-Wallace theorem as an ingredient in our subjective strategy for assigning probabilities on the basis of quantum states. I take it that this is in fact the subjectivist (or ‘pragmatist’) position on Everett proposed by Price (2010).

Thus, in all these cases (with the exception of Bohm of course), we have a view in which the probabilities of quantum mechanics are coherently understood as both purely subjective and ontic, which is what we set out to establish in this section.

5 Epistemic objective probabilities

We now need to discuss the case of epistemic objective probabilities. We shall do so in the context of Lewis’s treatment of objective probabilities, which however immediately raises an obstacle of principle. Recall that in Section 3 we identified epistemic probabilities as situated in the gap between what is knowable in principle and what is actually known, and that we settled (at least provisionally) on taking what is knowable in principle at time t to be the history of the world H_t up to that time. Now take our credence at t in some proposition A (about some events at time later than t). As mentioned, $\text{cr}_t(A) = \text{cr}(A|E_t \& T)$, where E_t are the facts at times up to t that we actually know at that time, and T are our theoretical assumptions. As in the case of (8) above, $\text{cr}_t(A)$ has thus the general form

$$\text{cr}(A|E_t \& T) = \sum_i \text{cr}(A|H_t^i \& T) \text{cr}(H_t^i|E_t \& T) . \quad (9)$$

Here, unless the theoretical assumptions T describe the situation as deterministic, the probabilities $\text{cr}(A|H_t^i \& T)$ will be non-trivial *ontic* probabilities. (In fact, it is only now that this statement is unproblematic, after we have argued in the last section that ontic probabilities can be equally subjective

or objective.²⁴) The epistemic component of our credence is given by the various $\text{cr}(H_t^i|E_t \& T)$. But if these probabilities are non-trivial, then according to Lewis they cannot be objective (i.e. Lewisian chances). Indeed, $\text{cr}(H_t^i|E_t \& T)$ is just our credence $\text{cr}_t(H_t^i)$ at t in the proposition H_t^i . But that is a proposition about past history, and according to Lewis chances of past events are always 0 or 1. Thus, non-trivial epistemic probabilities $\text{cr}(H_t^i|E_t \& T)$ cannot be objective.

We see that, given the PP, epistemic objective probabilities would seem to be an oxymoron: if we are ignorant of certain matters of fact, and if propositions about these facts are admissible in the sense of Lewis, then our non-trivial credences about these facts are not objective, because chances of past events are 0 or 1. It appears that Lewis is ruling out by definition that epistemic probabilities could ever be objective. As seemed to be the case with Popper, also for Lewis epistemic probabilities must be subjective.

For the sake of arguing that epistemic objective probabilities do make sense after all, we shall of course grant the existence of objective probabilities in the first place (*pace de Finetti*). Thus, assume that e.g. a quantum measurement of electron spin displays objective probabilities, say 50-50 chances of ‘up’ or ‘down’. We shall take these as objective, in the Lewisian sense that if I ask you to bet on the outcome of the next measurement, then conditionally on the chances it is rationally compelling for you to bet at equal odds.

Now consider the case in which I perform such a measurement but *do not show you the outcome yet*, and I ask you to bet. I make the following two claims:

- Your (subjective) probabilities about the outcome are epistemic: there is an outcome out there which is in principle knowable, and you employ probabilities only because you do not in fact know it.
- You are still going to bet at equal odds, and in fact this is *just as rationally compelling* as betting at those odds before the measurement.

The first claim is uncontroversial. The second one – your epistemic proba-

²⁴Indeed, we defined such probabilities as ontic, so unless we accept subjective ontic probabilities, we are committed to regarding these probabilities either as objective or (more generally!) as objectively wrong.

bilities after the measurement are rationally compelling, granted the probabilities before the measurement are – I take to be incontrovertible. From a Bayesian perspective, for instance, you have no new evidence. So it would be irrational to change your subjective probabilities after I have performed the experiment. But what is more, I shall now argue that your non-trivial probabilities are objective *for the same reason* as in Lewis’s discussion. That is, the rational justification for your probability assignment is still completely in the spirit of the PP, even though it does not respect its letter.

Indeed, recall what the PP (3) says, as an implicit definition of objective probabilities: iff there is a proposition X such that, for all choices of reasonable initial credence function and admissible propositions,

$$\text{cr}(A \mid X \ \& \ E_t \ \& \ T) = x \ , \tag{10}$$

then A has an objective probability and its value is x . Crucially, Lewis sees chances as functions of time. Thus he requires that our credence function at t should not incorporate information that is inadmissible *at the time t* . This corresponds to a betting situation in which you are allowed to know any matters of fact up to time t , and the chances fix the objectively fair odds for betting in this situation.

But in the case we are considering now, *the betting situation is different*. Indeed, you are not allowed to know what has happened after I have performed the measurement (even if the measurement was at a time $t' < t$). If in this betting situation you knew the outcome of the measurement, *that* now would be cheating. In this situation, the proposition that the chance at t is equal to x (in particular 0 or 1) is *itself* inadmissible information, and it is so for the same reason that Lewis holds a proposition about the future chance of an event to be inadmissible.

Indeed, in his version of the PP, Lewis is considering a betting context restricted to what we can in principle know at t (i.e. the past history H_t), and he asks for the best we can do in setting our rational credence in a future event A . The answer is that the best we can do is to set our credence to the chance of A at that time t . Information about how the chances evolve after t is inadmissible because it illicitly imports knowledge exceeding what we are allowed to know at t . If we consider restricting our betting context further so that what we are allowed to know is what we could in principle

have known at an earlier time t' , and now ask for the best we can do in setting our rational credence in an event at t , then information about the chance at t of that event has become inadmissible, again precisely because it exceeds the assumed restriction on what we are allowed to know. The best we can do is to set our credence to the chance of A at the earlier time t' .

The point is perhaps even clearer if we take Lewis's reformulated version (5) of the PP :

$$\text{ch}(A|H_t) = \text{cr}(A|H_t \& T_C) . \quad (11)$$

This tells us that if at t we include all admissible information, namely H_t and T_C , this information already implies what the chances are at t , so that the best we can do in terms of setting our credence at t , is to take any reasonable initial credence function and conditionalise on H_t and T_C .

But now if we assume that our epistemic context E_t is restricted to $H_{t'}$, clearly the best we can do is to take the information $H_{t'}$ we actually have and feed it into the complete theory of chance T_C , yielding as credence $\text{cr}(A|E_t \& T_C) = \text{ch}(A|H_{t'})$. *This* credence will now be rationally compelling, even if it is not equal to $\text{ch}(A|H_t)$.

The intuition is that in order to judge whether a subjective probability conditional on the epistemic context $E_t = H_{t'}$ is in fact objective, we need to compare it not to the objective probability conditional on the history H_t at the time of evaluation of our credence, but to the objective probability conditional on the history $H_{t'}$ that *matches* our epistemic context. And this now indeed provides the room for non-trivial objective probabilities even for events that could in principle be known.

6 Redrawing the subjective-objective distinction

We have argued above that one can make good sense both of subjective (indeed, purely subjective) ontic probabilities and of epistemic objective probabilities, thus arguing in favour of two unusual conceptual possibilities. But in other respects we had to be quite conservative. Indeed, the arguments above would have been weakened had we not tried to keep as much as possible to

the accepted use of the terms ‘epistemic’, ‘ontic’, ‘subjective’ and ‘objective’. Once we have seen the usefulness of admitting subjective ontic probabilities and epistemic objective probabilities, however, we might be interested in trying to redraw the subjective-objective distinction and the epistemic-ontic distinction in ways that are both more systematic and cleaner in the light of our new insights.

Let us start with the subjective-objective distinction. What we have just argued for in Section 5 is that there are more cases of rationally compelling credences than Lewis in fact allows for.²⁵

Lewis’s intuition is that chances are functions of *time*, and in the formulation of the PP this effectively forces him to privilege the maximal epistemic context H_t comprising all that is knowable in principle at t , or equivalently to *quantify* over all epistemic contexts E_t comprising what may be actually known at t . Thus, according to Lewis, we need to ask whether propositions exist that would fix our credences across all epistemic contexts admissible at time t .

But we have just argued that we can also ask whether propositions exist that would fix our credences when we are in a *particular* one of these epistemic contexts. Specifically, we considered the situation in which at t we only know all matters of fact up to $t' < t$. In that case, if we assume Lewisian chances exist, we know in fact that there are such propositions, namely the chances at t' .

Now we can ask that same question *in general*: do propositions exist that would fix our credences when we are in an *arbitrary* epistemic context E (e. g. any ever so partial specification of the history up to a time t)? If so, we shall talk of the chances not *at the time* t , but the chances *given the context* E . Lewisian chances at t are then just the special case when E has the form H_t .

Writing it down in formulas, we have a natural generalisation of the PP and of objective probabilities: given a context E , iff there is a proposition X such that, for all choices of reasonable initial credence function and any further

²⁵Many thanks to Ronnie Hermens for noting that this point was not spelled out clearly enough in a previous draft.

theoretical assumptions T (compatible with X),

$$\text{cr}(A \mid X \& E \& T) = x , \quad (12)$$

then A has an objective probability given the context E and its value is x . The weakest such proposition will be denoted $\text{ch}_E(A) = x$.

Now, similarly as with Lewis's own reformulation (5) of the PP, note that a complete theory of chance will now include also history-to-chance conditionals for any antecedent E for which chances given E exist. Therefore, not only will the following be an instance of the generalised PP (12),

$$\text{cr}(A \mid (\text{ch}_E(A) = x) \& E \& T_C) = x , \quad (13)$$

but we can further rewrite it as

$$\text{ch}_E(A) = \text{cr}(A \mid E \& T_C) . \quad (14)$$

Thus, also our generalised chances will have the kind of properties that Lewis derives for his own chances from his reformulation (5) of the original PP (in particular, they will obey the probability calculus).

Further, (14) makes vivid the intuition that if we have a complete and correct theory of (generalised) chance, the best we can rationally do if we are in the epistemic context E , is just to plug the information we *actually* have into the complete and correct theory of chance, and let that guide our expectations.

Finally, and very significantly, we see from (14) that chances in the more general sense are not functions of propositions and times, but functions of *pairs of propositions*, and they have the form of *conditional probabilities*. That is, in this picture *chances are fundamentally conditional*. I believe this has many advantages.²⁶

²⁶Note that also epistemic and ontic probabilities as discussed here are naturally seen as conditional (on epistemic and ontic contexts, respectively). That all probabilities should be fundamentally conditional has been championed by a number of authors (see e.g. Hájek (2011) and references therein). The idea in fact fits nicely with Lewis's own emphasis that already the notion of possibility is in fact to be naturally understood as a notion of *compossibility*, of what is possible keeping fixed a certain context. (Recall his lovely example in Lewis (1976) of whether he can speak Finnish as opposed to an ape, or cannot as opposed to a native speaker.)

From the point of view of the philosophy of physics, such a liberalisation of the PP means that all kinds of other probabilities used in physics can now be included in what could potentially be objective chances.

For instance, it allows one to consider the possibility of *high-level probabilities* being objective. By high-level probabilities I mean probabilities featuring in theories such as Boltzmannian statistical mechanics, which is a theory about *macrostates*, identified as coarse-grained descriptions of physical systems. Our generalised PP allows one among other things to apply a Lewisian-style analysis to probabilities conditional on such a macrostate of a physical system, thus introducing the possibility of objective chances even in case we are ignorant of the microstate of a system such as a gas (*pace* Popper).²⁷

Or take again the EPR setup: Bob performs a measurement, usually taken to collapse the state also of Alice's photon, so that the probabilities for her own measurement results are now $\cos^2(\alpha/2)$ and $\sin^2(\alpha/2)$. Let us grant that these probabilities are objective. In Section 3 we considered Alice's probability assignments in the epistemic context in which she does not know yet the result of the measurement on Bob's side, and treated them as partially epistemic. Now, however, we can argue that in Alice's epistemic context what is rationally compelling for her is to keep her probabilities unchanged, because we compare them with the quantum probabilities conditional on events only on *her* side.

Indeed, the situation is very much analogous to the simple example in which I perform a measurement but do not tell you the result. If anything, in the EPR example the case is even *more* compelling, because there is no way even in principle that Alice could know Bob's result if the two measurements are performed at spacelike separation. Thus, if we take probabilities to be encoded or determined by the quantum state, we have now made space for the idea that collapse does not act instantaneously at a distance, but along the *future light cone* (of Bob's measurement in this case).

²⁷And *pace* Lewis, who explicitly ruled out the possibility of such deterministic chances (1986, p. 118). Note that high-level deterministic chances relate to all kinds of interesting questions (e.g. emergence, the trade-off between predictive accuracy and explanatory power, the demarcation between laws and initial conditions, etc.). For recent literature on the debate about deterministic chances, see e.g. Glynn (2009), Lyon (2011), Emery (2013), and references therein.

Quite generally, it is clear that Lewis’s treatment of chances taken literally is incompatible with relativity, for Lewis’s chances as functions of time presuppose a notion of absolute simultaneity. Indeed, the idea that chance could be relativised to space-like hypersurfaces or more generally to circumstances which are independent of absolute time, e.g. properties of space-time regions, was proposed in Berkovitz (1998), who discusses this proposal with respect to causality and non-locality in the quantum realm.

A relativistic generalisation of Lewisian chances will automatically require more liberal contexts E – as we have proposed here – because there is no longer an absolute notion of simultaneity, or of ‘past’, or of ‘becoming’, determining which events we need to conditionalise upon to temporally evolve our chances. As is well known, in special relativity these notions can be modified in two ways: one either takes simultaneity to be relative to inertial frames or more generally to spacelike foliations, or one localises the notion of absolute past to the past light cone of a spacetime event (or spacelike segment or bounded spacetime region). Both indeed require generalising chances from just defining them relative to ‘past histories’ H_t to contexts that are the past of arbitrary spacelike hyperplanes or hypersurfaces, or, respectively, to contexts that are the past light cones of arbitrary events (or spacelike segments or bounded regions).

Conditionalising on the past of arbitrary hypersurfaces is what Myrvold (2000) has proposed as a way of conceptualising relativistic collapse in the case of quantum mechanical probabilities.²⁸ Conditionalising on past light cones instead is what Ismael (2011) has proposed as a way of conceptualising chances in the case of relativity. Applied to quantum mechanics, this corresponds to the EPR example as we have just revisited it, with a conception of collapse along the future light cone.²⁹

For yet another example, in recent work Adlam (2018a,b) has argued that

²⁸I actually prefer the 2000 archived version to the published paper (Myrvold 2002). More recently, Myrvold (2017a) has shown that any such theories of foliation-dependent collapse must lead to infinite energy production. I believe, however, that Myrvold’s idea that probabilities in special relativity (and in quantum field theory) can be defined conditional on events to the past of an arbitrary hypersurface can still be implemented (see Section 9 below).

²⁹I am indebted to both of these proposals and their authors in more ways than I can say. See also my Bacciagaluppi (2010a).

‘temporal locality’ and ‘objective chances’ risk becoming two mutually supporting dogmas of modern physics. Maybe slightly oversimplifying, the point is that the idea of objective chances defined at a time t and that of a physical state at time t being all that is needed to predict the physical state at times $t + \varepsilon$ in fact seem to naturally relate to and support each other. Adlam argues that this creates an obstacle to the development of e. g. non-Markovian, retrocausal, and ‘all-at-once’ theories.³⁰ But in fact, it is only the standard Lewisian form of objective chances *at time t* that may be inimical to such developments, and the present generalisation of the PP removes this particular obstacle.

For instance, suppose we have some physical theory involving non-Markovian probabilities. The Lewisian definition of chances actually does allow for non-Markovian chances, in the sense that $\text{ch}_t(A) = \text{ch}(A|H_t)$ now in general is a proposition that depends on the *whole* of the history H_t . However, if at t we only have information about some previous times t_1, \dots, t_n , the original PP has nothing to recommend us except to try and figure out the missing information, so as to be able to determine the chances $\text{ch}(A|H_t)$ at t . Our physical theory, instead, does contain also probabilities of the form $p(A|E_{t_1} \& \dots \& E_{t_n})$, where E_{t_i} describes the state of our physical system at time t_i . In order to be able to say that also these probabilities featuring in our non-Markovian theory are objective, we have to generalise chances to allow for contexts of the ‘gappy’ form $E = E_{t_1} \& \dots \& E_{t_n}$ in (12)–(14).

Note that $p(A|E_{t_1} \& \dots \& E_{t_n})$ is just a coarse-graining of probabilities of the form $p(A|H_t^j)$, where H_t^j is a possible history of the world up to t compatible with $E_{t_1} \& \dots \& E_{t_n}$, i.e. a coarse-graining of probabilities which in turn do allow a reading as Lewisian chances. In this sense, our generalisation (12) appears to be required by sheer consistency, if we want objective probabilities to be closed under coarse-grainings.³¹

³⁰Retrocausal theories (in particular in the context of quantum mechanics) have been vigorously championed by Price. All-at-once theories are ones in which probabilities for events are specified given general boundary conditions, rather than just initial (or final) conditions, and have been championed in particular by Wharton. See e.g. Price (1997), Wharton (2014) and the other references in Adlam (2018a). For further discussion see also Rijken (2018) (whom I thank for discussion of Adlam’s work).

³¹Of course also in the case of $E = H_{t'}$, the chance given E is the coarse-graining over all possible intervening histories between t' and t . But this particular coarse-graining does yield again a Lewisian chance, namely the chance at t' .

Finally, even leaving aside motivations from philosophy of physics, what may perhaps be decisive from the point of view of the philosophy of probability is that such a generalisation of the PP completely eliminates the need to restrict epistemic contexts to ‘admissible propositions’. The reason for worrying about admissibility was that chances were tacitly or explicitly thought to always refer to the context H_t , and the actual epistemic context had to be restricted to admissible propositions in order not to bring in any information not contained in H_t . By requiring that the (ontic) context of the chances exactly match the (epistemic) context of the credences, no such possibility of mismatch can arise, and no additional conditions of admissibility are needed.

I suspect Lewis himself might have welcomed such a simplification in the formulation of the PP, especially given his own qualified formulations when treating of admissible propositions in connection with the possibility of non-standard epistemic contexts in the presence of time travel, clairvoyance, and other unusual cases (as will in fact be the retrocausal and all-at-once theories considered by Adlam).

7 Redrawing the epistemic-ontic distinction

Generalising chances to arbitrary contexts E in our discussion of subjective and objective probabilities, suggests we might want to do something similar in our discussion of epistemic and ontic probabilities. While sufficient for the purpose of our examples so far, taking epistemic probabilities to be probabilities of past events (or any events determined by such events) is now too restrictive. As we did in the case of chances, we should not only want to be able to generalise to a relativistic setting, but also to be able to get rid of the restriction to past events altogether, allowing in fact for all kinds of non-standard epistemic contexts (clairvoyance, backwards-directed agents with memories of the future, omniscience, and what not), and allowing conversely for taking an ontic perspective also on probabilities that are about past events (e.g. in the case of coarse-graining).

This, however, also seems to raise a problem, namely that the distinction between epistemic and ontic probabilities becomes a purely pragmatic one at best. If we can extend or restrict at will which propositions we assume to

be knowable in principle, then any probability can be alternatively seen as epistemic or as ontic, and there is no substantive difference between the two.

To a large extent, we can bite the bullet, because the distinction between epistemic and ontic probabilities clearly does have such a pragmatic aspect. When in Section 5 we considered the example of the spin measurement with unknown outcome, in the given epistemic context those probabilities are of course epistemic. But if we are asked how we choose our probability assignments in that situation, we will say that we are simply applying the quantum probabilities that obtained at the time of the measurement, i. e. we are mentally *switching* to a context in which the indeterministic outcomes are still before us, and in which the probabilities predicted by quantum mechanics (which there we assumed are objective) are thus ontic, so that our theoretical assumptions about that ontic context can inform our probability assignments in our current epistemic context.

In so doing, however, we also seem to be performing a second switch, one that was in fact crucial to our arguing in Section 4 that ontic probabilities could also be subjective (e.g. when adopting a subjectivist position about probabilities of atomic decay). Indeed, while epistemic probabilities are about unknown matters of fact in the actual world, ontic probabilities as we used them in Section 4 are *theoretical* probabilities about alternative possibilities in our models.³²

We mentioned already in footnote 12 above a distinction between *material* propositions (propositions about matters of fact in the actual world) and *theoretical* propositions (propositions within our theoretical models about the world).³³ We should thus distinguish accordingly between different *dimensions of possibility*: material possibilities if we are wondering about what *is* the case in the material world, and theoretical possibilities if we are wondering about what *could be* or *might have been* the case.

³²And, indeed, already in Section 3 we suggested that ontic probabilities tend to be associated with indeterministic contexts, which we defined in terms of multiple possible futures in our theoretical models compatible with the present (or present and past) state of the world.

³³As mentioned above, one could of course distinguish further between the actual world and the material world (if one is a many-world theorist) or take some material propositions to be theoretical (if one is an empiricist), but for simplicity I neglect these possibilities.

In this sense, epistemic probabilities are clearly probabilities about material possibilities: this is the dimension of possibility in which beliefs live, and epistemic probabilities are measures over unknown matters of fact in the actual world (conditional on keeping the epistemic context fixed). But now we recognise an ambiguity in the standard use of ‘ontic probabilities’ as construed so far: we have been sliding between on the one hand taking also ontic probabilities to be *material* probabilities, namely measures over unknowable facts in the actual world (conditional on keeping the ontic context fixed, i.e. what is knowable in principle), and on the other hand taking ontic probabilities to be *theoretical* probabilities, namely measures over alternative possible worlds (now thinking of the fixed ontic context as what we are conditioning on *in the model*). In the radioactive decay example this means sliding between taking ontic probabilities as about what the actual facts in the future are, and taking them as measures over alternative possible worlds. Our examples of subjective ontic probabilities can be read in either way, and remain subjective under either reading, whether we identify ontic probabilities directly with the theoretical probabilities of our models (which in Section 4 we took to be purely subjective), or whether we take the ontic probabilities to be material, and *use* our subjective theoretical probabilities in fixing them.

It is not essential to choose between the two readings of ontic probabilities, as long as one can distinguish if required. Insofar as ontic probabilities feature in our expectations, one might prefer to think of them as material probabilities (since our expectations are about the actual world). On the other hand, our probabilistic statements about, say, future quantum events will tend to have modal force (in the sense of supporting causal and/or law-like statements), indicating that more often than not we intend such probabilities to be theoretical.³⁴

In this context, while in some of my examples I use the terminology of possible worlds, or refer to ontic probabilities as law-like, the important thing is indeed that our theoretical models and the probabilities that feature in them

³⁴Material probabilities by definition support only indicative conditionals: ‘If Beethoven was not born on 16th December, he was born on the 15th’ expresses our next best guess. Instead, when talking about atomic decay we more commonly use subjunctive or counterfactual conditionals, such as ‘Should this atom not decay in the next hour, it would have probability 0.5 of decaying during the following hour’.

have modal force, however understood. Our theoretical models may feature probabilities in terms of probabilistic laws, or in terms of probabilistic causes; if you object to possible worlds as a formalisation of modality (even though I take them in a very thin sense) or to thinking of probabilities as law-like, any alternative account of modal force will do.³⁵

Another point to note is that, while we have said that insofar as epistemic probabilities are material probabilities they do not have modal force, not all theoretical probabilities need in fact have modal force. Some probabilities in our models may be understood as contingent elements of a model. Such is typically the case for initial distributions, or more generally single-time distributions, even in theoretical models as, say, classical statistical mechanics. And, indeed, when discussing Popper’s worries in Section 1 we claimed as self-evident that the probability distributions in classical statistical mechanics are epistemic. But again we need to distinguish: as long as we are talking of our credences about what is in fact the case inside a bottle, these are indeed material and epistemic probabilities; as long, instead, as these are theoretical probability distributions over different alternative theoretical possibilities in our model, they do not range over epistemic possibilities. Rather, they appear to represent the contingent aspects of our theoretical model.³⁶ Thus we see that there is a certain ambiguity between material and theoretical probabilities also when we are talking about epistemic probabilities. Again, this ambiguity need not be worrying, as long as one is clearly aware of the material-theoretical distinction.³⁷

Indeed, to revisit one more time our standard example of mixed probabilities (unknown composition of a radioactive probe, known decay laws for the

³⁵My thanks to Niels van Miltenburg and to Albert Visser for pressing me on this point.

³⁶We shall refine this point in our discussion of time-symmetric theoretical models in Section 8.

³⁷It may, however, contribute to muddle the waters in the Popper and Jaynes examples, or indeed in modern discussions about ‘typicality’ in statistical mechanics: are we talking about the actual world (of which there is only one), or about possible worlds (of which there are many, at least in our models)? It is also not always clear which aspects of a model are indeed law-like and which are contingent. We shall see a class of examples in Section 8, but a much-discussed one would be whether or not the ‘Past Hypothesis’ can be thought of as a law – which it is in the neo-Humean ‘best systems’ approach (cf. Loewer 2007); thanks to Sean Gryb for raising this issue. For further discussion of typicality and/or the past hypothesis, see e.g. Goldstein (2001), Frigg and Werndl (2009), Wallace (2011), Pitowsky (2012), and Lazarovici and Reichert (2015).

various isotopes), we can now give multiple readings of these probabilities. We can see them as expressing our expectations about a material proposition in our future, which we evaluate by mixing our epistemic probabilities for the composition of the probe and our ontic probabilities (in the material sense) about what will happen in the various cases. In this case, we fix the values of the latter probabilities using the law-like probabilities of our theoretical models. Or we see them throughout as probabilities in a theoretical model, but with a contingent distribution over the isotopes in the radioactive probe. In that case, we can fix that distribution using our epistemic probabilities for the composition of the probe. Or yet again, we can think of such mixed probabilities as a *sui generis* mix of material and theoretical probabilities, which we however commonly resort to in practice.

How to use theoretical models (i.e. which probabilities to use) will depend pragmatically on our epistemic situation. Any given context can be seen as the epistemic context of some putative agent wondering about the actual world, or as an ontic context from which to judge the accessibility of other possible worlds. But in any *actual* epistemic context in which we wonder about what we do not know in the material world, we will happily and often resort to theoretical models to guide our expectations.

8 Remarks on time symmetry

Many more things could be said about applying the newly-drawn distinctions in practice. I shall limit myself in this section to a few remarks on the issue of time symmetry, and in the next one to some further remarks on the quantum state.

I mentioned in Section 3 that for the purpose of distinguishing epistemic and ontic probabilities I did not want to commit myself to a metaphysically thick notion of indeterminism, with a ‘fixed’ past and an ‘open’ future. Now we can see that such a notion is in fact incompatible with the way I have suggested drawing the epistemic-ontic distinction.

Indeed, as drawn in Section 7, the epistemic-ontic distinction allows for ontic probabilities also in the case we theoretically assume *backwards* indetermin-

ism (where models that coincide at a time t or for times $s \geq t$ need not coincide at times earlier than t). We can take the entire history to the *future* of t as an ontic context (call it \tilde{H}_t), and thus probabilities $p(A|\tilde{H}_t)$ about events *earlier* than t as ontic. This is plainly incompatible with tying ontic probabilities (whether in the material or the theoretical sense) to the idea of a fixed past and open future.

There are other reasons for rejecting the idea of fixed past and open future.³⁸ But the suggestion above seems to fly in the face of a number of arguments *for* a different metaphysical status of past and future themselves based on probabilistic considerations. It is these arguments I wish to address here, specifically by briefly recalling and reframing what I have already written in my Bacciagaluppi (2010b) against claims that past-to-future conditional probabilities will have a different status (namely ontic) from that of future-to-past conditional probabilities (allegedly epistemic). I refer to that paper for details and further references.

Of course we are all familiar with the notion that probabilistic systems can display strongly time-asymmetric behaviour. To fix the ideas, take atomic decay, modelled simplistically as a two-level classical Markov process (we shall discuss quantum mechanics proper in the next section), with an excited state and a de-excited state, a (comparatively large) probability α for decay from the excited state to the de-excited state in unit time, and a (comparatively small) probability ε for spontaneous re-excitation in unit time. We shall take the process to be homogeneous, i.e. the transition probabilities to be time-translation invariant.

Any initial probability distribution over the two states will converge in time to an equilibrium distribution in which most of the atoms are de-excited (a fraction $\frac{\alpha}{\alpha+\varepsilon}$), and only a few are excited (a fraction $\frac{\varepsilon}{\alpha+\varepsilon}$), and the transition rate from the excited to the de-excited state will asymptotically match the

³⁸In particular, I agree with Saunders (2002) that standard moves to save ‘relativistic becoming’ in fact undermine such metaphysically thick notions. Rather, I believe that the correct way of understanding the ‘openness’ of the future, the ‘flow’ of time and similar intuitions is via the logic of temporal perspectives as developed by Ismael (2017), which straightforwardly generalises to relativistic contexts, since it exactly captures the perspective of an IGUS (‘Information Gathering and Utilising System’) along a time-like worldline. For introducing me to the pleasures of time symmetry I would like to thank in particular Huw Price.

transition rate from the de-excited to the excited state (so-called ‘detailed balance’ condition).

This convergence to equilibrium is familiar time-directed behaviour. Of course, given the joint distributions for two times $s > t$, transition probabilities can be defined both from t to s and from s to t (by conditionalising on the state at t and at s , respectively), but the example can be used in a number of ways to argue that these transition probabilities forwards and backwards in time must have a different status. In the language of this paper:

(A) Were backwards probabilities also ontic and equal to the forwards probabilities, then the system would converge to equilibrium also towards the past, and the system would have to be stationary (Sober 1993).

(B) While by assumption the forwards transition probabilities are time-translation invariant, the backwards transition probabilities in general are not, suggesting that they have a different status (Arntzenius 1995).

(C) If in addition to the forwards transition probabilities also the backwards transition probabilities were law-like, the initial distribution would also inherit this law-like status, while it is clearly contingent – indeed can be freely chosen (Watanabe 1965).

The problem with all these arguments is that they take the observed backwards *frequencies* as estimates for the putative ontic backwards *probabilities*. But frequencies can be used as estimates for probabilities only when we have reasons to think our samples are unbiased.³⁹ For instance, we would never use post-selected ensembles to estimate forwards transition probabilities. However, we routinely use *pre*-selected ensembles, especially when we set up experiments ourselves and in fact freely choose the initial conditions. Thus, while we normally use such experiments to estimate forwards transition probabilities, we plainly cannot rely on them to estimate the backwards ones.

It turns out that this time-asymmetric behaviour (and a large class of more general examples) can indeed be modelled using stationary stochastic processes in which the transition probabilities are time-symmetric, and in which the time asymmetry is introduced by the choice of special initial (rather than

³⁹That this is not always licensed is clear e.g. from discussion of causal loops (Berkovitz 1998, 2001).

final) *frequencies*. Thus, in modelling such time-asymmetric behaviour there is no need to use time-asymmetric ontic probabilities or to deny ontic status altogether to the backwards transition probabilities of the process. The apparent asymmetry comes in because we generally have asymmetric *epistemic access* to a system: we tend to know aspects of its past history, and thus are able to pre-select, while its future history will generally be unknowable.⁴⁰

If one assumes that the transition probabilities are time-symmetric (or time-symmetric except for some overall drift⁴¹), this places severe constraints on the overall (multi-time) distributions that define the process. Not only is the process stationary if it is homogeneous, but the process may well be uniquely determined by its forwards and backwards transition probabilities.⁴² Thus, in particular, the entire process inherits the ontic status of the transition probabilities.

This will seem puzzling, if we are used to thinking of transition probabilities as ontic, and of initial distributions as epistemic, so that the probabilities of the process should have a mixed character (as in the examples we discussed beginning in Section 3). But as we saw in Section 7, we standardly align the distinction between epistemic and ontic probabilities with the distinction between contingent and law-like aspects of our theoretical probabilistic models. The *theoretical probabilities* of the process may very well be even uniquely determined by the forwards and backwards transition probabilities, but our

⁴⁰Of course, this strategy is analogous to that of modelling time-directed behaviour in the deterministic case using time-symmetric laws and special initial conditions. The same point was made independently and at the same time by Uffink (2007, Section 7), specifically in the context of probabilistic modelling of classical statistical mechanics. A further explicit example of time-symmetric ontic probabilities is provided by Nelson's (1966, 1985) stochastic mechanics, where the quantum mechanical wavefunction and the Schrödinger equation are derived from an underlying time-symmetric diffusion process. See also Bacciagaluppi (2005) and references therein. The notorious gap in Nelson's derivation pointed out by Wallstrom (1989) has now been convincingly filled through the modification of the theory recently proposed by Derakhshani (2017).

⁴¹In the case of Nelson's mechanics, for instance, one has a systematic drift velocity that equals the familiar de Broglie-Bohm velocity (which in fact also changes sign under time reversal). If one subtracts the systematic component, transition probabilities are then exactly time-symmetric.

⁴²More precisely, it will be uniquely determined within each 'ergodic class'. (Note that ergodicity results are far easier to prove for stochastic processes than for deterministic dynamical systems. See again the references in Bacciagaluppi (2010b) for details.)

theoretical models will still provide scope for contingent elements through the notion of *samples* of the process. We then feed our epistemic probabilities into the theoretical model via constraints on the frequencies in the samples, which will typically be initial constraints because of our time-asymmetric access to the empirical world.⁴³

9 Further remarks on the quantum state

We have just seen that, classically, a stochastic process may be nothing over and above an encoding of its forwards and backwards transition probabilities. I now wish to briefly hint at the possibility that the same may be true also in quantum mechanics.⁴⁴

In quantum mechanics, the role of determining transition probabilities belongs to the quantum state (either the fixed Heisenberg-picture state, or the totality of the time-evolving Schrödinger-picture or interaction-picture states⁴⁵). The quantum state, however, is traditionally seen as more than simply encoding transition probabilities, rather as a *physical object* in itself, which in particular determines these probabilities. One obvious exception is qBism, where however I have argued in Section 4 above that the rejection of the quantum state as a physical object is not sufficiently motivated.⁴⁶

⁴³Without going into details, these distinctions will be relevant to debates about the ‘neo-Boltzmannian’ notion of typicality (see the references in footnote 37), and about the notion of ‘sub-quantum disequilibrium’ in theories like pilot-wave theory or Nelsonian mechanics (for the former see e.g. Valentini 1996, for the latter see Bacciagaluppi 2012).

⁴⁴This section can be skipped without affecting the rest of the paper. The discussion is substantially influenced by Myrvold (2000), but should be taken as an elaboration rather than exposition of his views. See also Bacciagaluppi (2010a). For Myrvold’s own recent views on the subject, see Myrvold (2017b, 2019).

⁴⁵The interaction picture is obtained from the Schrödinger picture by absorbing the free evolution into the operators. In other words, one encodes the free terms in the evolution in the evolution of the operators, and encodes the interaction in the evolution of the state.

⁴⁶Another obvious exception is Nelson’s stochastic mechanics (mentioned in footnote 40), but there it is not the quantum mechanical transition probabilities that are fundamental but the transition probabilities of the underlying diffusion process. While also a number of Bohmians (cf. Goldstein and Zanghì 2013) think of the wavefunction not as a material entity but as nomological, as a codification of the law-like motion of the particles, there is no mathematical derivation in pilot-wave theory of the wavefunction

In this section I want to take a look at a different approach to quantum theory, where I believe there are strong constraints to taking the quantum state as an independently existing object, and reasons to take it indeed as just encoding (ontic, law-like) transition probabilities, largely independently of the issue of whether these probabilities should be subjective or objective.⁴⁷

This approach is the foliation-dependent approach to relativistic collapse. As is well-known, non-relativistic formulations of collapse will postulate that the quantum state collapses instantaneously at a distance (whether the collapse is spontaneous or the traditional collapse upon measurement). So, for instance, in an EPR scenario the collapse will depend on whether Alice's or Bob's measurement takes place first. However, *predictions* for the outcomes of the measurements are independent of their order because the measurements commute: we obtain the same probabilities for the same pairs of measurement results. This suggests that one might be able to obtain a relativistically invariant formulation of collapse after all. One of the main strategies for attempting this takes quantum states to be defined relative to spacelike hyperplanes or more generally relative to arbitrary spacelike hypersurfaces, and collapse accordingly to be defined relative to spacelike foliations. This option has been pioneered for a long time by Fleming (see e.g. 1986, 1989, 1996), and analysed conceptually in great detail by Myrvold (2000) (see also Bacciagaluppi (2010a)).⁴⁸

A convenient formal framework for foliation-dependent collapse is provided by the Tomonaga-Schwinger formalism of quantum field theory, in which the (interaction-picture) quantum state is explicitly hypersurface-dependent, and unitary evolution in the presence of interactions is represented as a local evolution of the quantum state from one hypersurface to the next (and is independent of the choice of foliation between two hypersurfaces S and S'). A non-unitary collapse evolution can be included in the same way, as proposed

and Schrödinger equation from more fundamental entities comparable to that in stochastic mechanics. The wavefunction can be thought of as derivative only if one takes a quite radical relationist view, like the one proposed by Esfeld (2017).

⁴⁷For a sketch of how I think this relates to the PBR theorem (Pusey, Barrett and Rudolph 2012), see below.

⁴⁸The other main strategy is to consider collapse to be occurring along light cones (either along the future light cone or along the past light cone, as in the proposal by Hellwig and Kraus (1970), and it arguably allows one to retain the picture of quantum states as physical objects (Bacciagaluppi 2010a). No such theory has been developed in detail, however.

explicitly e.g. by Nicosini and Rimini (2003), provided one can ensure that the evolution is also independent of the foliation chosen.

The substantial difference between the unitary and non-unitary case is that in the unitary case, the local state on a segment of a spacelike hypersurface (say, the reduced state of Alice’s electron) is independent of whether the segment, say, $T \cap T'$ is considered to be a segment of a hypersurface T or of a different hypersurface T' . In the non-unitary case, in general the restriction $\psi_{T|T \cap T'}$ on $T \cap T'$ of a state ψ_T on T is different from the restriction $\psi_{T'|T \cap T'}$ on $T \cap T'$ of the state $\psi_{T'}$ on T' . Indeed, in the EPR case ψ_T could be the singlet state, so that $\psi_{T|T \cap T'}$ is maximally mixed, while T' could be to the future of Bob’s measurement, so that $\psi_{T'|T \cap T'}$ is, say, the state $|+_y\rangle$. In particular, whether Alice’s electron is entangled with Bob’s will depend on the hypersurface one considers (or the frame of reference – Myrvold (2000) aptly calls this the ‘relativity of entanglement’).

This now raises issues of how to interpret such a collapse theory, specifically in a way that could give rise to the manifest image of localised objects in space and time. The standard options for doing so go under the names of ‘wavefunction ontology’, ‘mass-density ontology’, and ‘flash ontology’, depending on whether one takes the physical objects of the theory to be the quantum states, local mass densities defined via the states, or the collapse events themselves. Hypersurface-dependence of quantum states appears to rule out the wavefunction ontology, precisely because if quantum states are hypersurface-dependent, they do not have unique local restrictions, so an ontology based on them fails to underpin any local reality. Similarly, mass-density ontology is ruled out, because hypersurface-dependent states do not define uniquely any local mass densities. The only option that could give rise to a localised manifest image of reality appears to be the flash ontology. This means that hypersurface-dependent states indeed merely encode *probabilities for physical events to the future of a spacelike hypersurface, conditional on physical events to the past of that hypersurface*. If so, as Myrvold (2000, footnote 7) remarks, it is unsatisfactory that we should not yet have a formulation of the theory directly in terms of transition probabilities, but require an additional theoretical object, the quantum state, to determine these probabilities. (As we saw in Section 6, also Adlam (2018a,b) is critical of thinking of a separate physical object as determining the quantum probabilities.)

What I wish to suggest is that, as in the classical case of Section 8, the use of a separate quantum state to calculate transition probabilities may be an artefact of our own limited epistemic perspective, and that one needs no more than forwards and backwards transition probabilities corresponding to a ‘stationary process’. The form of this process is easy to guess, reasoning by time symmetry from the fact that the future boundary condition we standardly use is the maximally mixed state. If we take that to be the case because we happen to lack information about the future, and therefore do not post-select, that suggests that if we want to remove the pre-selection that we normally do, and thus get to the correct theoretical probabilities, we need to take the maximally mixed state also as the past boundary condition. The appearance of a contingent quantum state would then follow from the fact that we have epistemic access to past collapses, and conditionalise upon them when making predictions, in general mixing our epistemic probabilities about such past collapses (‘results of measurements’) and ontic transition probabilities provided by the theory.

Of course in general the maximally mixed state is unnormalisable, but *conditional* probabilities will be well-defined. In such a ‘no-state’ quantum theory not only are conditional probabilities fundamental, but unconditional probabilities in general do not even exist.⁴⁹

No-state foliation-dependent flash-ontology quantum field theory would thus be a way of directly implementing Myrvold’s idea that conditional probabilities should be determined by the theory directly, and not by way of the quantum state. Better predictions will be obtained the more events one conditionalises upon stretching back into the past. The theory is thus essentially

⁴⁹Rather related ideas are common in the decoherent histories or consistent histories literature. Griffiths (1984) does not use an initial and a final state, but an initial and a final projection of the same kind as the other projections in a history, so his consistent histories formalism is in fact a no-state quantum theory. Hartle (1998) suggests that the fundamental formula for the decoherence functional is the two-state one, but that we are well-shielded from the future boundary condition so that two-state quantum theory is predictively equivalent to standard one-state quantum theory (i.e. we can assume the maximally mixed state as future boundary condition). If in one-state quantum theory we assume that we are similarly shielded from the initial boundary condition, then by analogy one-state quantum theory is in fact predictively equivalent to no-state quantum theory (i.e. we can assume the maximally mixed state also as past boundary condition). Thus, it could just as well be that no-state quantum theory is fundamental. (Of course we need to include pre-selection on the known preparation history.)

non-Markovian (thus temporally non-local, in line with Adlam’s views). As mentioned in footnote 28, Myrvold (2017a), developing remarks by Pearle (1993), has shown that foliation-dependent collapse theories always run into a problem of infinite energy production, but I believe a no-state theory in fact escapes this conclusion.⁵⁰

Such a proposal has in fact been put forward in the non-relativistic case by Bedingham and Maroney (2017), in a paper in which they discuss in general the possibility of time-symmetric collapse theories. As with the requirement of Lorentz invariance, the requirement of time symmetry makes it problematic to think of the collapsing quantum state as a physical object, because the sequence of collapsing states defined by the collapse equations applied in the past-to-future direction is not the time reversal of the sequence of states defined by the collapse equations applied in the future-to-past direction. That is, the quantum state is time-direction-dependent. However, under very mild conditions, from the application of the equations in the two directions of time one does obtain the same probabilities for the same collapse events. This suggests, again as in the relativistic case, that a time-symmetric theory of collapse requires a flash ontology. And, as I am suggesting in the case of hypersurface-dependence, Bedingham and Maroney suggest one might want to get rid altogether of the quantum state (since the collapse mechanism suggests it is asymptotically maximally mixed).⁵¹

⁵⁰The idea behind Myrvold’s proof is extremely simple. Leaving technicalities aside, because of the Lorentz-invariance of the vacuum state, infinite energy production would arise if localised collapse operators did not leave the vacuum invariant on average. Thus, one should require that local collapse operators leave the vacuum invariant. But since a local operator commutes with local operators at spacelike separation, and since (by the Reeh-Schlieder theorem) applications of the latter can approximate any global state, it follows that in order to avoid infinite energy production collapse operators must leave every state invariant. Thus there is no collapse.

One way out that is already present in the literature (and which motivates Myrvold’s theorem) is to tie the collapse to some ‘non-standard’ fields (which commute with themselves at spacelike and timelike separation), as done by both Bedingham (2011) and Pearle (2015) himself. Another possibility, favoured by Myrvold (2018a), is to interpret the result as supporting the idea that collapse is tied to gravity: in the curvature-free case of Minkowski spacetime, the theorem shows there is no collapse. As mentioned, however, I believe that a no-state proposal will also circumvent the problem of infinite energy production, because collapse operators will automatically leave the (also Lorentz-invariant) maximally mixed state invariant on average.

⁵¹I should briefly remark on how I think these suggestions relate to the the work of

One can find an analogy for the approach discussed in this section by looking at Heisenberg’s views on ‘quantum mechanics’ (i.e. matrix mechanics) as expressed in the 1920s in opposition to Schrödinger’s views on wave mechanics.⁵² The evidence suggests that Heisenberg did not believe in the existence of the quantum state as a physical object, or of collapse as a physical process. Rather, the original physical picture behind matrix mechanics was one in which quantum systems performed quantum jumps between stationary states. This then evolved into a picture of quantum systems performing stochastic transitions between values of measured quantities. Transition probabilities for Heisenberg were ‘out there’ in the world (i.e. presumably objective as well as ontic). If a quantity had been measured, then it would perform transitions according to the quantum probabilities. The observed statistics would conform to the usual rules of the probability calculus, in particular the law of total probability. But if a quantity had not been measured, there would be nothing to perform the transitions, giving rise to the phenomenon of ‘interference of probabilities’.

Schrödinger’s quantum states were for Heisenberg nothing but a convenient tool for calculating (forwards) transition probabilities, and any state that led to the correct probabilities could be used. (Note that, classically, forwards transition probabilities are independent of the chosen initial conditions.) Hence in particular, the movability of the ‘Heisenberg cut’ and Heisenberg’s cavalier attitude to collapse,⁵³ and the peculiar mix of ‘subjective’ and ‘ob-

Pusey, Barrett and Rudolph (2012), who famously draw an epistemic-ontic distinction for the quantum state, and who conclude from their PBR theorem that the quantum state is ontic. An ontic state for PBR is an independently existing object that determines the ontic probabilities (in my sense) for results of measurements. And an epistemic reading of the quantum state for PBR is a reading of it as an epistemic probability distribution (in my sense) over the ontic states of a system. Thus, the PBR theorem seems to push for a wavefunction ontology. In a no-state proposal (as here or in Bedingham and Maroney (2017)), however, the probabilities for future (or past) collapse events are fully determined by the set of past (or future) collapse events. The contingent quantum state (in fact whether defined along hypersurfaces or along lightcones) can thus be *identified* with such a set of events, which is indeed ontic, thus removing the apparent contradiction. Whether this is the correct way of thinking about this issue of course deserves further scrutiny.

⁵²For more detailed accounts, see Bacciagaluppi and Valentini (2009), Bacciagaluppi (2008), Bacciagaluppi and Crull (2009), and Bacciagaluppi, Crull and Maroney (2017).

⁵³For an example of (Born and) Heisenberg using a very unusual choice of ‘collapsed states’ to calculate quantum probabilities, see Bacciagaluppi and Valentini (2009, Sect. 6.1.2).

jective' elements in talking about quantum probabilities. And now, finally, despite the shrouding in 'Copenhagen mist', we can perhaps start to make sense of some of Heisenberg's intuitions as expressed in 'The Copenhagen interpretation of quantum theory' (Heisenberg 1958, Chap. 3).

10 Hume and de Finetti

This paper has suggested a framework for thinking of probabilities in a way that distinguishes between subjective and epistemic probabilities. Two strategies have been used: liberalising the contexts with respect to which one can define various kinds of probabilities, and making a liberal use of probabilities over theoretical possibilities. I wish to conclude by putting the latter into a wider perspective and very briefly sketching the kind of subjectivist and anti-metaphysical position that I wish to endorse.⁵⁴

In forming theoretical models, whether they involve causal, probabilistic or any other form of reasoning, we aim at describing the behaviour of our target systems in a way that guides our expectations. To do so, our models need to have modal force, to describe law-like rather than accidental behaviour. It is up for grabs what this modal force is. We can have an anti-Humean view of causes, laws or probabilities, as some kind of necessary connections, but Hume (setting aside the finer scruples of Hume scholarship) famously denied the existence of such necessary connections. This traditionally gives rise to the problem of induction. Hume himself proposed a 'sceptical solution': we are predisposed to come to consider certain connections between events as causal through a process of habituation, which then forms the basis of our very sophisticated reasonings about matters of fact ('Elasticity, gravity, cohesion of parts, communication of motion by impulse; these are probably the ultimate causes and principles which we shall ever discover in nature; and we may esteem ourselves sufficiently happy, if, by accurate enquiry and reasoning, we can trace up the particular phenomena to, or near to, these general principles', Hume 1748, IV.I.26). There is no guarantee that our

⁵⁴To this position, I would wish to add a measure of empiricism (Bacciagaluppi 2019). For a similar view, see Brown and Ben Porath (this volume). On these matters (and several others dealt with in this paper) I am earnestly indebted to many discussions with Jenann Ismael over the years. Any confusions are entirely my own, however.

expectations will come out true, but these forms of reasoning have served us well.

There is no need, however, for Humeans to restrict themselves to thinking of theoretical models in terms of Hume's own analysis of causation. Any theoretical models are indeed patterns we project onto the world of matters of fact in order to reason about it. Modern-day pragmatists about causation are essentially following Hume, even though they substitute manipulation for habituation (Menzies and Price 1993). And the same is true of subjectivists about probability. Subjective probabilities are patterns that we form in our head, and impose on matters of fact to try and order them for our pragmatic purposes.

Self-styled latter-day Humeans often attempt to give objective underpinnings to the laws and probabilities in our models. The best systems approach takes Humean laws to be the axioms in the best possible systematisation of the totality of events, and Lewis's 'Subjectivist Guide' is a manifesto for doing so in the case of subjective probabilities.

But the Humean may well have no need of such objective underpinnings. Radical subjectivists about probability have been saying so for a long time. There is no sense in which probabilistic judgements are right or wrong, just as there is no necessary connection between causes and effects. By and large through a process similar to habituation (Bayesian conditioning!) we form the idea of probabilistic connections, which then forms the basis of our further reasonings about matters of fact. There is no guarantee that our expectations will come out true, but these forms of reasoning have served us well.

For the subjectivist, or pragmatist, such are all our theoretical models of science, including quantum mechanics. And we judge them in terms of the standard criteria of science: first and foremost the empirical criterion of past performance, as well as other criteria such as simplicity, expectation of fruitfulness, and so on. In this sense, there are no objective chances in the sense of strict Lewisian rationality, no objective laws nor objective modality in either an anti-Humean or even neo-Humean sense, but merely in the sense of the pragmatic rationality of science.

Indeed, the best systems analysis provides a plausible if schematic picture of our inductive practices, but there is no point to applying it at the 'end

of time', when we no longer need to project any predicates. Any laws that we formulate along the way, any notions of modality that we construe are only the result of analysing the epistemically available data, and as such are only projections of what has been successful so far. From a Humean point of view, this seems to me as it should be, just as Hume himself was content with his sceptical solution to the problem of induction.

And, as a matter of fact, in the relatively neglected Section VI of the Enquiry, Hume explicitly discusses probabilities in terms of forming expectations, again through habit, about what will happen in certain proportions of cases. Thus Hume was perhaps the first de Finettian...

Acknowledgement

I am deeply grateful to Meir Hemmo and Orly Shenker for the invitation to contribute to this wonderful volume, as well as for their patience. I have given thanks along the way, both for matters of detail and for inspiration and stimuli stretching back many years, but I must add my gratitude to students at Aberdeen, Utrecht and the 5th Tübingen Summer School in the History and Philosophy of Science, where I used bits of this material for teaching, and to the audience at two Philosophy of Science seminars at Utrecht, where I presented earlier versions of this paper, in particular to Sean Gryb, Niels van Miltenburg, Wim Mol, Albert Visser and Nick Wiggershaus. Special thanks go to Ronnie Hermens for precious comments on the first written draft, further thanks to Harvey Brown, Alan Hájek and Jenann Ismael for comments and encouraging remarks on the final draft, and very special thanks to Jossi Berkovitz for a close reading of the final draft and many detailed suggestions and comments.

References

Adlam, E. (2018a): Spooky action at a temporal distance. *Entropy*, 20(1), 41–60

Adlam, E. (2018b): A tale of two anachronisms. Plenary talk given at Foundations 2018, Utrecht University, 13 July 2018. <https://foundations2018.sites.uu.nl/>

Arntzenius, F. (1995): Indeterminism and the direction of time. *Topoi* 14, 67–81

Bacciagaluppi, G. (2005): A conceptual introduction to Nelsons mechanics. In: Buccheri, R., Saniga, M., and Avshalom, E. (eds.) *Endophysics, Time, Quantum and the Subjective*, pp. 367–388. World Scientific, Singapore. Revised version at <http://philsci-archive.pitt.edu/8853/>

Bacciagaluppi, G. (2008): The statistical interpretation according to Born and Heisenberg. In: Joas, C., Lehner, C., and Renn, J. (eds.) *HQ-1: Conference on the History of Quantum Physics*. MPIWG Preprint Series, vol. 350, pp. 269–288. MPIWG, Berlin. <http://www.mpiwg-berlin.mpg.de/en/resources/preprints.html>

Bacciagaluppi, G. (2010a): Collapse theories as beable theories. *Manuscripto* 33(1), 19–54. <http://philsci-archive.pitt.edu/8876/>

Bacciagaluppi, G. (2010b): Probability and time symmetry in classical Markov processes. In: Suárez, M. (ed.) *Probabilities, Causes and Propensities in Physics*. Synthese Library, vol. 347, pp. 41–60. Springer, Dordrecht. <http://philsci-archive.pitt.edu/archive/00003534/>

Bacciagaluppi, G. (2012): Non-equilibrium in stochastic mechanics. *Journal of Physics: Conference Series* 361(1), article 012017 <http://philsci-archive.pitt.edu/9120/>

Bacciagaluppi, G. (2014): A critic looks at qBism. In: Galavotti, M.C., Dieks, D., Gonzalez, W., Hartmann, S., Uebel, T., and Weber, M. (eds.) *New Directions in the Philosophy of Science*, pp. 403–416. Springer, Cham. <http://philsci-archive.pitt.edu/9803/>

Bacciagaluppi, G. (2016): Quantum probability – an introduction. In: Hájek, A., and Hitchcock, C. (eds.) *The Oxford Handbook of Probability and Philosophy*, pp. 545–572. Oxford, OUP. <http://philsci-archive.pitt.edu/10614/>

- Bacciagaluppi, G. (2019): Adaptive empiricism. In: D'Ariano, G.M., Robbiati Bianchi, A., and Veca, S. (eds.) *Lost in Physics and Metaphysics – Questioni di Realismo Scientifico*, pp. 99–113. Istituto Lombardo Accademia di Scienze e Lettere, Milano. <https://doi.org/10.4081/incontri.2019.465>
- Bacciagaluppi, G., and Crull, E. (2009): Heisenberg (and Schrödinger, and Pauli) on hidden variables. *Studies in History and Philosophy of Modern Physics* 40, 374–382. <http://philsci-archive.pitt.edu/archive/00004759/>
- Bacciagaluppi, G., Crull, E., and Maroney, O. (2017): Jordan's derivation of blackbody fluctuations. *Studies in History and Philosophy of Modern Physics* 60, 23–34. <http://philsci-archive.pitt.edu/13021/>
- Bacciagaluppi, G., and Valentini, A. (2009): *Quantum Theory at the Crossroads – Reconsidering the 1927 Solvay Conference*. CUP, Cambridge
- Bedingham, D. (2011): Relativistic state reduction dynamics. *Foundations of Physics* 41, 686–704
- Bedingham, D., and Maroney, O. (2017): Time reversal symmetry and collapse models. *Foundations of Physics* 47, 670–696
- Bell, J.S. (1984): Beables for quantum field theory. Preprint CERN-TH-4035. Reprinted in: *Speakable and Unspeakable in Quantum Mechanics*, pp. 173–180. CUP, Cambridge (1987)
- Berkovitz, J. (1998): Aspects of quantum non-locality II – superluminal causation and relativity. *Studies in History and Philosophy of Modern Physics* 29(4), 509–545.
- Berkovitz, J. (2001): On chance in causal loops. *Mind* 110(437), 1–23
- Berkovitz, J. (2012): The world according to de Finetti – on de Finetti's theory of probability and its application to quantum mechanics. In: Ben-Menahem, Y., and Hemmo, M. (eds.) *Probability in Physics*, pp. 249–280. Springer, Heidelberg
- Berkovitz, J. (2019): On de Finetti's instrumentalist philosophy of probabil-

ity. *European Journal for Philosophy of Science* 9, article 25

Brown, H. (forthcoming): The reality of the wavefunction – old arguments and new. In: Cordero, A. (ed.) *Philosophers Look at Quantum Mechanics*. Synthese Library, vol. 406. Springer (due September 2019)

Brown, H., and Ben Porath, G. (this volume): Everettian probabilities, the Deutsch-Wallace theorem and the Principal Principle

Bub, J. (2007): Quantum probabilities as degrees of belief. *Studies in History and Philosophy of Modern Physics* 38, 232–254

Bub, J. (2010): Quantum probabilities – an information-theoretic interpretation. In: Hartmann, S., and Beisbart, C. (eds.) *Probabilities in Physics*, pp. 231–262. OUP, Oxford

Bub, J., and Pitowsky, I. (1985): Critical notice – Sir Karl R. Popper, Postscript to the Logic of Scientific Discovery. *Canadian Journal of Philosophy* 15(3), 539–552

de Finetti, B. (1970): *Theory of Probability*. Wiley, New York

Derakhshani, M. (2017): Stochastic mechanics without ad hoc quantization – theory and applications to semiclassical gravity. Doctoral dissertation, Utrecht University. <https://arxiv.org/pdf/1804.01394.pdf>

Diaconis, P., Holmes, S., and Montgomery, R. (2007): Dynamical bias in the coin toss. *SIAM Review* 49(2), 211–235

Emery, N. (2013): Chance, possibility, and explanation. *The British Journal for the Philosophy of Science* 66(1), 95–120

Esfeld, M. (2017): A proposal for a minimalist ontology. *Synthese*, 1–17

Fleming, G. (1986): On a Lorentz invariant quantum theory of measurement. In: Greenberger, D.M. (ed.) *New Techniques and Ideas in Quantum Measurement Theory*. *Annals of the New York Academy of Sciences*, vol. 480, 574–575. New York Academy of Science, New York

Fleming, G. (1989): Lorentz invariant state reduction, and localization. In: Fine, A., and Forbes, M. (eds.) *PSA 1988*, vol. 2, pp. 112–126. Philosophy

of Science Association, East Lansing, MI

Fleming, G. (1996): Just how radical is hyperplane dependence? In: Clifton, R. (ed.) *Perspectives on Quantum Reality – Non-Relativistic, Relativistic, and Field-Theoretic*, pp. 11–28. Kluwer Academic, Dordrecht

Frigg, R., and Hoefer, C. (2007): Probability in GRW theory. *Studies in History and Philosophy of Modern Physics* 38(2), 371–389

Frigg, R., and Hoefer, C. (2010): Determinism and chance from a Humean perspective. In: Dieks, D., Gonzalez, W., Hartmann, S., Weber, M., Stadler, F., and Uebel, T. (eds.) *The Present Situation in the Philosophy of Science*, pp. 351–371. Springer, Berlin

Frigg, R., and Werndl, C. (2012): Demystifying typicality. *Philosophy of Science* 79(5), 917–929

Fuchs, C.A. (2002): Quantum mechanics as quantum information (and only a little more). arXiv preprint [quant-ph/0205039](https://arxiv.org/abs/quant-ph/0205039)

Fuchs, C.A. (2014): Introducing qBism. In Galavotti, M.C., Dieks, D., Gonzalez, W., Hartmann, S., Uebel, T., and Weber, M. (eds.) *New Directions in the Philosophy of Science*, pp. 385–402. Springer, Cham

Glynn, L. (2009): Deterministic chance. *The British Journal for the Philosophy of Science* 61(1), 51–80

Goldstein, S. (2001): Boltzmann’s approach to statistical mechanics. In: Bricmont, J., Dürr, D., Galavotti, M.C., Ghirardi, G.C., Petruccione, F., and Zanghì, N. (eds.) *Chance in Physics. Lecture Notes in Physics*, vol. 574, pp. 39–54. Springer, Berlin

Goldstein, S., and Zanghì, N. (2013): Reality and the role of the wave function in quantum theory. In: Ney, A., and Albert, D.Z. (eds.) *The Wave Function – Essays on the Metaphysics of Quantum Mechanics*, pp. 91–109. OUP, Oxford

Greaves, H., and Myrvold, W. (2010): Everett and Evidence. In: Saunders, S., Barrett, J., Kent, A., and Wallace, D. (eds.) *Many Worlds? Everett, Quantum Theory, and Reality*, pp. 264–304. OUP, Oxford

- Griffiths, R. B. (1984): Consistent histories and the interpretation of quantum mechanics. *Journal of Statistical Physics* 36(1–2), 219–272
- Hájek, A. (2011): Conditional probability. In: Gabbay, D.M., Thagard, P., Woods, J., Bandyopadhyay, P.S., and Forster, M.R. (eds.) *Philosophy of Statistics. Handbook of the Philosophy of Science*, vol. 7, pp. 99–135. North-Holland, Dordrecht
- Hartle, J. B. (1998): Quantum pasts and the utility of history. *Physica Scripta T* 76, 67–77
- Heisenberg, W. (1958): *Physics and Philosophy – The Revolution in Modern Science*. Harper, New York
- Hellwig, K.-E., and Kraus, K. (1970): Formal description of measurements in local quantum field theory. *Physical Review D* 1, 566–571
- Howard, D. (2004): Who invented the Copenhagen interpretation? A study in mythology. *Philosophy of Science* 71(5), 669–682
- Hume, D. (1748): *An Enquiry Concerning Human Understanding*. A. Millar, London
- Ismael, J. (2011): A modest proposal about chance. *The Journal of Philosophy* 108(8), 416–442
- Ismael, J. (2017): Passage, flow, and the logic of temporal perspectives. In: Bouton, C., and Huneman, P. (eds.) *Time of Nature and the Nature of Time. Boston Studies in the Philosophy and History of Science*, vol. 326, pp. 23–38. Springer, Cham
- Jaynes, E.T. (1985): Some random observations. *Synthese* 63(1), 115–138
- Lazarovici, D., and Reichert, P. (2015): Typicality, irreversibility and the status of macroscopic laws. *Erkenntnis* 80(4), 689–716
- Lewis, D. (1976): The paradoxes of time travel. *American Philosophical Quarterly* 13(2), 145–152
- Lewis D. (1980): A subjectivist’s guide to objective chance. In: Jeffrey, R.C. (ed.) *Studies in Inductive Logic and Probability*, vol. 2, 263–293. Reprinted

- in: *Philosophical Papers*, vol. 2, pp. 83–113. OUP, Oxford (1986)
- Lewis, D. (1986): Postscripts to ‘A subjectivist’s guide to objective chance’. In: *Philosophical Papers*, vol. 2, pp. 114–132. OUP, Oxford
- Loewer, B. (2004): David Lewis’s Humean theory of objective chance. *Philosophy of Science* 71, 1115–1125
- Loewer, B. (2007): Counterfactuals and the second law. In: Price, H., and Corry, R. (eds.) *Causation, Physics, and the Constitution of Reality – Russell’s Republic Revisited*, pp. 293–326. OUP, Oxford
- Lyon, A. (2011): Deterministic probability – neither chance nor credence. *Synthese* 182(3), 413–432
- Menzies, P., and Price, H. (1993): Causation as a secondary quality. *British Journal for the Philosophy of Science* 44, 187–203
- Myrvold, W. (2000): Einstein’s untimely burial. *PhilSci* preprint <http://philsci-archive.pitt.edu/222/>
- Myrvold, W. (2002): On peaceful coexistence – is the collapse postulate incompatible with relativity? *Studies in History and Philosophy of Modern Physics* 33(3), 435–466
- Myrvold, W. (2017a): Relativistic Markovian dynamical collapse theories must employ nonstandard degrees of freedom. *Physical Review A* 96(6), 062116
- Myrvold, W. (2017b): Ontology for collapse theories. In: Shan Gao (ed.) *Collapse of the Wave Function – Models, Ontology, Origin, and Implications*, pp. 99–126. CUP, Cambridge
- Myrvold, W. (2018a): Private communication, New Directions in Philosophy of Physics, Viterbo, June 2018. <http://carnap.umd.edu/philphysics/newdirections18.html>
- Myrvold, W. (2018b): Philosophical issues in quantum theory. In: Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition). <https://plato.stanford.edu/archives/fall2018/entries/qt-issues/>

- Myrvold, W. (2019): Ontology for relativistic collapse theories. In: Lombardi, O., Fortin, S., López, C., and Holik, F. (eds.) *Quantum Worlds*, pp. 9–31. CUP, Cambridge. <http://philsci-archive.pitt.edu/14718/>
- Myrvold, W. (this volume): Subjectivists about quantum probabilities should be realists about quantum states
- Nelson, E. (1966): Derivation of the Schrödinger equation from Newtonian mechanics. *Physical Review* 150, 1079–1085
- Nelson, E. (1985): *Quantum Fluctuations*. Princeton University Press, Princeton
- Nicosini, O., and Rimini, A. (2003): Relativistic spontaneous localization – a proposal. *Foundations of Physics* 33(7), 1061–1084
- Pearle, P. (1993): Ways to describe dynamical state-vector reduction. *Physical Review A* 48(2), 913–923
- Pearle, P. (2015): Relativistic dynamical collapse model. *Physical Review D* 91, article 105012
- Pitowsky, I. (1989): *Quantum Probability, Quantum Logic*. Lecture Notes in Physics, vol. 321. Springer, Berlin
- Pitowsky, I. (1994): George Boole’s ‘Conditions of Possible Experience’ and the Quantum Puzzle. *British Journal for the Philosophy of Science* 45, 95–125
- Pitowsky, I. (2003): Betting on the outcomes of measurements – a Bayesian theory of quantum probability. *Studies in History and Philosophy of Modern Physics* 34(3), 395–414
- Pitowsky, I. (2007): Quantum mechanics as a theory of probability. In: Demopoulos, W., and Pitowsky, I. (eds.) *Physical Theory and its Interpretation – Festschrift in Honor of Jeffrey Bub*. Western Ontario Series in Philosophy of Science, vol. 72, pp. 213–240. Springer, New York
- Pitowsky, I. (2012): Typicality and the role of the Lebesgue measure in statistical mechanics. In: Ben-Menahem, Y., and Hemmo, M. (eds.) *Probability in Physics*, pp. 41–58. Springer, Heidelberg

- Popper, K.R. (1982): *Quantum Theory and the Schism in Physics. Postscript to the Logic of Scientific Discovery*, vol. III. Bartley III, W.W. (ed.). Hutchinson, London
- Price, H. (1997): *Time's Arrow and Archimedes' Point – New Directions for the Physics of Time*. OUP, Oxford
- Price, H. (2010): Decisions, decisions, decisions – can Savage salvage Everettian probability? In: Saunders, S., Barrett, J., Kent, A., and Wallace, D. (eds.) *Many Worlds? Everett, Quantum Theory, and Reality*, pp. 369–391. OUP, Oxford
- Pusey, M. F., Barrett, J., and Rudolph, T. (2012): On the reality of the quantum state. *Nature Physics* 8(6), 475–478
- Rijken, S. (2018): *Spacelike and Timelike Non-Locality*. Masters Thesis, History and Philosophy of Science, Utrecht University. <https://dspace.library.uu.nl/handle/1874/376363>
- Saunders, S. (2002): How relativity contradicts presentism. *Royal Institute of Philosophy Supplements* 50, 277–292
- Saunders, S., Barrett, J., Kent, A., and Wallace, D. (eds.) (2010): *Many Worlds? Everett, Quantum Theory, and Reality*. OUP, Oxford
- Sober, E. (1993): Temporally oriented laws. *Synthese* 94, 171–189
- Uffink, J. (2007): Compendium of the foundations of classical statistical physics. In: Butterfield, J., and Earman, J. (eds.) *Handbook of the Philosophy of Physics*, part B, pp. 923–1074. North-Holland, Amsterdam
- Valentini, A. (1996): Pilot-wave theory of fields, gravitation and cosmology. In: Cushing, J.T., Fine, A., and Goldstein, S. (eds.) *Bohmian Mechanics and Quantum Theory – An Appraisal*, pp. 45–66. Springer, Dordrecht
- Wallstrom, T.C. (1989): On the derivation of the Schrödinger equation from stochastic mechanics. *Foundations of Physics Letters* 2(2), 113–126
- Watanabe, S. (1965): Conditional probability in physics. *Progress of Theoretical Physics Supplement*, extra issue, 135–167

Wharton, K. (2014): Quantum states as ordinary information. *Information* 5, 190–208

Wallace, D. (2010): How to prove the Born rule. In: Saunders, S., Barrett, J., Kent, A., and Wallace, D. (eds.) *Many Worlds? Everett, Quantum Theory, and Reality*, pp. 227–263. OUP, Oxford

Wallace, D. (2011): The Logic of the past hypothesis. PhilSci preprint <http://philsci-archive.pitt.edu/8894/>