NOÛS

NOÛS 45:2 (2011) 269–299

Deterministic Chance

ANTONY EAGLE
University of Oxford

Abstract

I sketch a new constraint on chance, which connects chance ascriptions closely with ascriptions of ability, and more specifically with 'CAN'-claims. This connection between chance and ability has some claim to be a platitude; moreover, it exposes the debate over deterministic chance to the extensive literature on (in)compatibilism about free will. The upshot is that a prima facie case for the tenability of deterministic chance can be made. But the main thrust of the paper is to draw attention to the connection between the truth conditions of sentences involving 'CAN' and 'CHANCE', and argue for the context sensitivity of each term. Awareness of this context sensitivity has consequences for the evaluation of particular philosophical arguments for (in)compatibilism when they are presented in particular contexts.

The consensus amongst philosophers is probably that non-trivial objective chances require indeterminism. One finds little argument for this view; even those who trouble to state the view explicitly (Lewis, 1980: 120) don't go to the further effort of justifying it, perhaps because it might seem obvious. After all, under determinism, wouldn't someone who knew the past and the laws know all there was to know about the future? So what could our use of probability be under determinism but a side-effect of our ignorance?

Against this silent consensus, there has been a small but vocal contingent of philosophers who have attempted to defend the possibility of deterministic chance. The main motivation has been to understand the impeccably objective role of probabilities in classical statistical mechanics, a deterministic theory for all intents and purposes in which probability has an irreducible role. In a variety of ways, Clark (1987), Glynn (2010), Hoefer (2007), Ismael (2009), Sober (2010), and perhaps most influentially Loewer (2001), have all tried to defend the idea that objective chance and determinism can happily coexist.

Unfortunately, I don't believe any of these accounts succeed. In this I am in agreement with Jonathan Schaffer, who has recently concluded that 'there cannot be a deterministic chance on any of the [common] conceptions, since the platitudes about the role of chance do not allow it' (Schaffer, 2007: 136). I won't stop here to develop or explain Schaffer's argument; suffice it to say that he thinks all the extant accounts of deterministic chance, and natural generalisations of them, all fail to capture one or another central truth about chance. With this view I concur.

But I'm not willing to follow Schaffer all the way to his pessimistic conclusion that *no* account of deterministic chance is possible. I think there is a further central truth about chance, reflection on which suggests that deterministic chance, in some sense, remains an open possibility. Given this possibility, concluding that chance is incompatible with determinism is at least premature. The thesis I'm arguing for is that there are some theories which are deterministic, and under which not all true chance ascriptions need be trivial. So I'm not arguing that any such theory is actually true; nor do I argue that *every* deterministic theory has genuine chances. Nevertheless, I will argue that some theories of independent interest, and which have been thought to trivialise chance—such as classical physics—do not in fact do so without additional assumptions. It may be for all I say that, given the way actual physics will turn out, there are only trivial chances. But if so, I argue, that will not be due solely to actual physics being deterministic, but will be due to actual physics (and how we talk about it) meeting further stronger conditions in addition to, or even instead of, determinism. It will become clearer in the sequel what kinds of conditions might require that the only true chance ascriptions be trivial ones (§4).

However, my main conclusion (in section 6) is more general. The main importance that the possibility or otherwise of deterministic chance has is not its intrinsic interest, significant though that may be. It is in the philosophical consequences that it has been argued the thesis has: for the understanding of chance, ability, determinism, and the status of non-fundamental science. And here I think my discussion helps clarify to a considerable degree precisely what does and does not follow from typical incompatibilist claims about chance and determinism, and thus contributes to a better understanding of these other central issues in the metaphysics of science. To foreshadow: I will argue there is a subtle but crucial context sensitivity in many sentences used by (in)compatibilists to express their views. But that is to come. First, I want to motivate my neglected central truth about chance.

1. Ability and Chance

I and my informants, all competent speakers of English, agree that the following two claims about some ordinary coin, while clearly meaningful, are obviously false:[1]

(1) a. This coin can land heads when tossed; still, it has no chance of landing heads when tossed.
b. This coin can't land heads when tossed; still, it has some chance of landing heads when tossed.

I don't suppose that the evident badness of the sentences in (1) will convince everyone. Perhaps some will attempt complicated pragmatic explanations for the contradictory appearance of these claims. I will leave that to those so inclined; the more reasonable opinion, as far as I can see, is to take the badness of these claims at face value. These sentences are bad because they are false; and they are false, not because of any *a posteriori* knowledge we have about the coin in question, but because, in each claim, the truth of one conjunct excludes the truth of the other.

The simplest explanation for this fact is that ascriptions of non-zero chance entail, and are entailed by, 'CAN' claims. These 'CAN' claims are themselves most plausibly understood in the *dynamic* sense, as objective ascriptions of ability.[2] While 'CAN', like other English modal auxiliaries, also has both deontic and epistemic readings (though the epistemic reading generally occurs only under negation), those readings are clearly not capable of explaining the cases in (1) (Palmer, 2001: 9–10, 76–8). In drawing this connection between chance ascriptions and abilities I do not mean to appeal to any 'thick' sense of ability—in the sense I'm concerned with, something has an ability just in case a 'CAN' claim is true of it. Having an ability, in the sense in question, will not be a phenomenon peculiar to agents, let alone agents with libertarian free will.

Any satisfactory account of chance must explain the connection between chance and ability manifest in (1). There are a number of ways of doing this. Levi is explicit in arguing that 'attributions of chances presuppose attributions of dispositions and abilities' (1990: 131), and uses this claim in the course of articulating a tendency interpretation of chance propensities. Similar in this respect is Peirce (1910: 169), who argues that chances are connected with the 'habits' of dice and other chance set ups, as evidenced by long term frequencies of behaviour.[3] For myself, I'd like just to draw attention to the fact that the sentences in (1) aren't just 'CAN' claims—they embed the ability ascriptions in a *habitual* construction. A habitual generalisation is a sentence of the form 'X ϕs WHEN C' (Krifka *et al*., 1995); I think the lesson to be drawn is that chance claims are connected not to bare occasional abilities, but to ascriptions of habitual (or generic) abilities to the chance setup.[4] But the points I'm going to make don't depend sensitively on the exact details.

So I'm going to set the details of the precise implementation of the chance-ability connection aside, and concentrate instead on a simple version of that connection for illustrative purposes. I don't think the principle I'll propose is precisely correct as it stands, neglecting as it does the habitual constructions in (1). But the ways in which it is incorrect don't matter for present purposes,

for if this simple principle can be shown to allow for deterministic chance, it can be argued that more sophisticated principles descended from it will do so as well, and those principles are the basis for serious candidate accounts of chance.

The principle I refer to is this:

> **The Chance-Ability Principle (CAP)** Where ⌜X⌝ is a noun phrase, and ⌜ϕ⌝ a complement verb phrase, the chance of X ϕ-ing exceeds zero iff X can ϕ. (Similarly, as 'MUST' is dual to 'CAN', the chance of X ϕ-ing is one iff X must ϕ.)

CAP is the simplest explanation of the falsity of the sentences in (1), and while it is somewhat naive, it still captures something very important about chance. CAP obviously entails that each of the sentences in (1) are contradictory, which in turn successfully predicts our intuitive reaction to those sentences.[5]

*

Obviously, CAP doesn't provide a theory of chance; it is at best able to distinguish outcomes with trivial chances from those without, and leaves unconstrained the assignment of more fine-grained numerical chances to those propositions. By itself, it yields only a constraint on 'qualitative' chance, distinguishing only between trivial and non-trivial chances.[6] CAP, or something very much like it, is true (and known to be true), and any correct theory of chance must include it, whatever turns out to ground its truth. But CAP doesn't place many constraints on what that theory of chance will look like. It is incompatible with actual frequentism, because many things that can happen do not, and so have non-zero chance while they have zero actual frequency. The truth of CAP thus is a further argument against frequency theories of chance. But CAP is consistent with our best current views about chance, including propensity theories (as mentioned above), as well as sophisticated empiricist theories like Lewis' best system account (Lewis, 1994; Loewer, 2004).[7]

The CAP does face an obvious problem. Many readers will have recalled that standard probability theory permits possible outcomes to be assigned zero probability—that in some cases (uniform distributions over continuous spaces) it even demands that there be possible outcomes with zero probability. For example, the only chance that can be consistently uniformly assigned to a point-sized dart hitting any point on a continuous dartboard is zero; yet all of those outcomes are possible. It might look like CAP clashes with this obvious fact about the probability calculus. Yet this clash follows only on the assumption that these two claims are true:

(2) 'CAN' in CAP is synonymous with 'POSSIBLY', and
(3) 'CHANCE' in CAP designates a quantity that must obey the standard probability calculus in its treatment of probability zero outcomes.

Yet there is in fact not much reason to suppose that either of these claims must be true. I think the claim that should go is (2), as I'll now argue; I relegate the discussion of (3) to Appendix A.

The current linguistic consensus is overwhelmingly that dynamic (ability-attributing) uses of 'CAN' should be distinguished from the bare alethic modal 'IT IS POSSIBLE THAT' (Kratzer, 1977; Palmer, 2001). I don't contend that 'POSSIBLY' is always alethic, invariantly expressing logical or even metaphysical possibility. The flexibility of English modals is notorious, and we often mean 'CAN' by 'POSSIBLY'. Still, 'POSSIBLY' does sometimes express quite unrestricted notions of possibility (particularly when stressed in an utterance), and certainly less restricted than dynamic 'CAN'. This means that some events which are logically and perhaps metaphysically possible nevertheless cannot happen.[8] This leaves it open that all the putative counterexamples to CAP (and its more sophisticated descendants) involve events which, while possible, cannot come to pass. This seems intuitively correct to me in all the cases normally introduced to demonstrate possible zero-chance outcomes, like the point-sized dart case—since no one *can make* a point-sized dart, I see no way to claim that an event involving such a dart *can happen*, though such events certainly strike me as possible. Just because something is possible doesn't mean anyone or anything is able to bring it about, and as such (2) is not plausibly true.[9]

Moreover, it is very plausible that probability zero events cannot come to pass. Suppose we had a theory T, which assigned probability zero to some event, and we observed that event come to pass. That would be much stronger evidence that T was false than that an extraordinarily unlikely event had occurred. (It would, on a Bayesian model, reduce our posterior credence in T to zero.) No chance zero events ever actually come to pass, since were they to do so, they must not have really had zero chance (we just thought that they did). While there will be a possible world w where some event occurs which is assigned zero chance by the correct chance theory of the actual world, that event will not be assigned zero chance by the correct chance theory of w.[10] Thus it is plausible that our CAP is compatible with the standard probability calculus.

A further line of objection attempts to block both of these responses. The objection is, even if point sized darts can't be made, still, we have the point-sized *centre* of any finite sized dart's tip, and (given standard models of spacetime) there will be a point where the centre lands. Similarly, for any temporally extended event e there is the unextended event, the beginning of e, and this will occur at some point of time. So, even if ordinary events don't have probability zero, events which are part of ordinary events can have probability zero; and since ordinary events can happen, those events which are part of them can happen. This objection seems to me a strong one; my inclination is to respond by denying that counterexamples based on this kind of idealisation, like the existence of a point-sized centre of a

dart or of a precise beginning of an event, are compelling and cogent. It is only according to quite contentious philosophical theorising that whenever an ordinary dart hits a dartboard, the distinct event of the centre of the dart hitting a point on the dartboard also occurs. If matter is anything like how we think it to be (that is, not perfectly homogenous), there will not be a part of the dart coming into contact with a part of the dartboard to correspond to this second extraordinary event. So it is perfectly in accordance with both commonsense and a straightforward understanding of science that we deny the existence of these further idealisations. No problem arises for the probability calculus, since a probability function defined over a sigma field of ordinary events (like a dart's hitting a dartboard) needn't be defined over these extraordinary constituent events (like the centre of the dart hitting a point on the dartboard). So I'm inclined to think that, despite appearances, this last sort of case is in the same position as the first sort—the idealisation of a point-sized centre of a dart is just as much an idealisation as the dart itself being point-sized, and that idealisation is equally undeserving of our serious credence. Probability spaces which include uncountably many outcomes, each of which is possible but has probability zero, are often useful in modelling real world phenomena. The mathematics of such spaces is often easier to handle than trying to handle these phenomena in a discrete space. But that doesn't mean that the phenomena modelled—the outcomes which really, actually, bear the chances—share every feature with the idealisation. It is mathematically nice to model ordinary objects as if they had point centres of mass; but ordinary objects are not points. So too with chances.[11]

I'm aware that these brief responses won't satisfy all objectors. I concede that this objection, and others I will mention below, remain live. The aim of the present paper isn't to give a full defence of CAP, and I've already conceded that we should not ultimately endorse that principle as formulated (so the real question is whether the above cases are counterexamples to a refined version of CAP not yet on the table, a question I see no way to answer in advance). My aim, rather, is to draw attention to and discuss the neglected platitude about the connection between chance and ability, which may potentially refocus the discussion about deterministic chance. The prima facie case given here for the tenability of CAP, and *a fortiori* for the more developed proposals based upon it, suffices to defuse the urgency of the present objection, and allows me to pursue the main goal of the paper with justified optimism that some defensible principle along the present lines can be found.

2. Ability and the Chance Role

Before I discuss how CAP interacts with compatibilism about chance, I need to argue that theories which endorse CAP can be genuine theories of chance. Schaffer's main argument against existing accounts of deterministic chance is

that they cannot fulfil the chance role; it would be nice to know that CAP is at least consistent with the rest of what we know about chance. Summarising a lot of recent work in the philosophy of probability, Schaffer provides a list of principles philosophers have taken to characterise the chance role (2007: §4), and I propose here to show that CAP is consistent with these other things we know about chance, and that collectively therefore they can serve to fix the referent of 'chance', as the thing which best fits the principles collectively, including CAP.

The other principles, which I will discuss in turn below, are these: (i) the Realization Principle; (ii) the Basic Chance Principle; (iii) the Principal Principle; (iv) the Future Principle; (v) the Intrinsicness Principle; (vi) the Causal Transition Constraint; and (vii) the Lawful Magnitude Principle. To this list, I am arguing, we should add CAP. There may well be still further natural and intuitive principles about chance, but Schaffer's list provides a convenient summary of the most popular claims philosophers have taken to collectively define the chance role.

As well as listing these principles about chance, Schaffer gives formal statements that he thinks articulate the substantial content of those principles. While I accept that these seven principles are more or less obviously true, I don't think Schaffer's formal versions are at all obvious, and so I reject his way of spelling out the content of the principles. Indeed, incompatibilism is an almost irresistible outcome of articulating the formal principles in the way Schaffer does. I'm therefore going to show that the connection between chance and ability can fit together with, and in many cases provide support for, the basic principles, not Schaffer's formal articulations.

For example, Schaffer states what he calls the *Realization Principle*: that if the chance of p at t in w is non-zero, there exists a world w' which matches w in history and laws up to t and where p is true. This formal principle is motivated by the informal claim that 'if there is a non-zero chance of p, this should entail that p is possible, and indeed that p is compossible with the circumstances' (Schaffer, 2007: 124). While I have to dissent from the formal Realization Principle—it quite clearly leads to incompatibilism in the present framework—I see no reason why that should bar me from agreeing with the intuitive motivation.[12] Of course I would represent that intuition slightly differently. Indeed I would argue that it is plausibly motivated by CAP itself, in the form: if there is a non-zero chance of p then p can happen. The ability ascription using 'CAN' is obviously similar to the modal 'POSSIBLE' that appears in the informal statement of the intuition. I rather suspect in fact that the content of the intuition is indeterminate as to whether it involves a pure alethic modality as opposed to an ability modal; it is not plausible that the folk intuition really has a precise content that excludes CAP. So I conclude that, in this case, CAP can accommodate the intuition.

And as it goes for the Realization Principle, so it goes for Schaffer's other formal claims: CAP is at least consistent with the principles (though on

occasion they need to be rendered using a slightly different formal claim) and in some cases it provides another buttress for the principle in question. Most obviously, CAP is consistent with the *Basic Chance Principle*. Again, set aside the formal version of that principle.[13] The argument that motivates the BCP is this:

> [I]f the chance of A is positive there must be a possible future in which A is true.... But what kinds of worlds can have futures that ground the fact that there is a positive present chance of A in the actual world? Not just any old worlds. The future occurrence of A in some world quite unlike the actual world in its past or present states of affairs is quite irrelevant to the actual present chance of A. (Bigelow *et al.*, 1993: 459)

A very similar argument to that given above for the RP shows that CAP is consistent with the BCP. For X's present ability to ϕ must be grounded in a possibility that X does ϕ, as Bigelow et al. argue. Moreover, it cannot be grounded in just any possibility, because (as I've just argued) 'CAN' is not equivalent to a bare alethic possibility. Rather, the present ability is grounded in the present nature and circumstances of X being consistent with X ϕ-ing, so the chance-grounding possibility must be one where X ϕs under circumstances largely similar to those X is currently to be found in. And this, in turn, supports the BCP, which entails that if p has some chance, then p comes to pass in a possibility very much like the current situation.

Another principle can be extracted from the obvious fact that our knowledge of X's abilities, the things X can do, should constrain our credences in various propositions about what X *will* do. We should assign no credence to X ϕ-ing iff X can't ϕ, and we know it can't. This, in conjunction with CAP, entails the obvious truth that we should have positive credence in X ϕ-ing just in case X has some chance of ϕ-ing. In this way information about chances constrains rational credence, in line with Lewis' *Principal Principle* (1980). The Principal Principle, recall, is that in the absence of definite information about the outcome of some chance event, one's conditional credence in that event—conditional on the chance of that event being x—should (rationally) also be equal to x. Or, more briefly, without oracular information, our beliefs about chances should constrain our credence in the events with those chances.

The CAP features in an interesting argument for the Principal Principle. For this norm connecting credence and ability seems right:

> **Credence-If-Can** One should have some (non-zero) credence in E's coming to pass if one believes that E can come to pass.

(Whether we do so in fact is another matter.) In conjunction with CAP, the Credence-If-Can principle entails that one should assign some non-zero

credence to something's coming to pass if one believes it has some non-zero chance of coming to pass. This norm entails that one's unconditional credence in an outcome should be non-zero if one believes the chance of that outcome is non-zero. It is plausible to move from this last claim, to the principle that one's conditional credence in an outcome, conditional on the chance of that outcome being non-zero, should also be non-zero. And this latter norm on reasonable credence is the special case of the Principal Principle for 'has some chance'. The Principal Principle in full generality may emerge from a more sophisticated version of CAP; but at least the start of an argument for the Principal Principle comes from the version of CAP floated here.

In fact, I think adopting something like CAP is the only way to consistently maintain the Principal Principle. If CAP did not hold, there could be cases where an outcome could happen, and reasonable agents therefore assign non-zero credence to that outcome in line with the Credence-If-Can principle, even while the theory of chance assigned zero probability to that outcome. (If there could be no such cases, CAP would hold.) But then the Principal Principle would fail, since reasonable agents would assign non-zero credence conditional on belief in zero chance. The only way out would be to say that information about what can happen is inadmissible, which seems implausible. So I conclude that, assuming the highly plausible Credence-Can principle, the only views of chance which can vindicate the Principal Principle are those which also vindicate CAP (or something like it).

Schaffer's fourth principle is the *Future Principle*, that the only non-trivial chances concern the future. This principle derives its plausibility from the truth that 'what's past is no longer chancy' (Lewis, 1980: 93). But since what is past can't normally be affected, if X can ϕ, and it's not the case that X must ϕ (so that X has a non-trivial chance of ϕ-ing), then X's ϕ-ing must be a future possibility (as must X's failing to ϕ). This suffices to capture the FP, at least under the plausible assumption that there is an asymmetry between the past and future with respect to what we are now able to affect.[14]

The fifth principle, the *Intrinsicness Principle* demonstrates clearly that Schaffer's formal principles are no mere restatements of intuition, involving as it does in his formulation ungainly reference to mereological sums of events. What the principle is intended to capture is that 'chances should remain constant across intrinsically duplicate trials' (Schaffer, 2007: 125). Our standard practices involving ability ascriptions certainly respect this principle: we are inclined to assign the same abilities to objects which are the same,when they are in the same circumstances (both of which should be held constant in an intrinsic duplicate trial of X: it wouldn't suffice to duplicate only the coin and not the method of tossing if one wanted to duplicate a coin toss). Insofar as abilities depend on these occurrent facts, they respect the Intrinsicness Principle. This feature is respected by standard accounts of the dynamic ability modal 'CAN' (as discussed below in §5), because 'in

all cases, the modal base [over which 'CAN' quantifies] includes propositions expressing the subject's stable intrinsic qualities' (Portner, 2009: 201–2). And since one has to have an ability at the time one exercises it, if there is a chancy transition between events *c* and *e*, then the chance must have a non-zero value between the time of *c* and the time of *e*, just as required by Schaffer's *Causal Transition Constraint*.

The seventh and final principle is the *Lawful Magnitude Principle* (see also Lange 2006). Informally, this requires that chances should assume the values assigned them in the laws of nature. On the face of it, this poses no problem for CAP. If the laws explicitly mention chances, then as the laws are true, the chances will be as the laws say they are. CAP then requires that the entities involved have a lawful ability to bring that chancy outcome about, which is obviously right. From another direction: if the fundamental physical theory is indeterministic, certain abilities follow, primarily generated by the possibility of various future outcomes in conjunction with the RP and BCP, which then entails (by CAP) non-trivial chances.[15]

Not all laws explicitly mention chances or abilities. The laws may simply be silent on the values of chances. There may still be chances, just as there are human beings despite the laws of our physics being apparently silent concerning our existence. But one might wonder: is there ever a case where the laws don't mention chances or abilities, and where that silence should be interpreted as a commitment to the denial of the existence of abilities?

One case which naturally suggests itself, and which is the focus of the remainder of this paper, is the case of deterministic laws. I'll go on to argue that this case doesn't arise, and the existence of chances and abilities is compatible with the laws being deterministic (or more cautiously, compatibilism about chance and determinism is at least no worse off than compatibilism about ability and determinism). So I don't regard this as ultimately a problematic case for ability-involving accounts of chance.

Are there other problem cases? Generally one might suspect these cases could arise if a strong variety of reductionism were true. The variety I am imagining would insist that when some laws, in conjunction with history, suffice for a complete account of every event occurring at a world, the only the things and properties explicitly mentioned in the laws exist at that world. Thus stated, this view would suffice to show that under determinism, non-trivial abilities may not exist, since they are not absolutely required for the account of occurrent events. But this reductionist view is entirely implausible: it doesn't seem that many, if any, modal facts would be true (since one needn't say how things possibly or necessarily are to say how they are). Unsurprising then that facts about abilities turn out false too.

This reductionism is implausible in another relevant way too. Some English sentences are context sensitive, so that the very same sentence can express a different proposition on different occasions of use. So 'AUSTRALIA IS BIG' can express different propositions depending on the contextually

determined threshold for how large something must be to count as big. Some of the propositions expressed might be ⟨AUSTRALIA IS LARGER THAN 10 KM^2⟩ (true) or ⟨AUSTRALIA IS LARGER THAN RUSSIA⟩ (false). Suppose, plausibly, that every proposition thereby expressed follows from the laws and history. The reductionism just sketched would presumably deny that there is any property of *bigness*, but to claim on that basis that natural language sentences involving 'BIG' are defective would be incorrect. If, as I shall argue below (following Kratzer (1977) and Lewis (1979)), ability ('CAN') claims and chance ascriptions are context-sensitive, then the silence of the laws concerning them is no reason at all to suppose that there are no abilities (chances) grounded by true ability (chance) ascriptions.[16]

I conclude, therefore, that CAP is consistent with, and indeed forms a very natural package with, the rest of what we ordinarily take ourselves to know about chances, and that it is a constraint that needs to be met by any legitimate contender to be chance. We may therefore rely on it when discussing what follows from the interaction between chance and determinism.

3. Chance, Ability and Determinism

Having shown that CAP can be included in an account of genuine chances, and may well be a mandatory feature of any such account, what does it tell us about the relationship between chance and determinism? In short, nothing much yet: without an account of 'CAN', the principle yields no substantive conclusions about chance and determinism.

But it does enable us to connect the debate over determinism and chance with the ongoing and extensive discussions over the compatibility of determinism and free action. For ability ascriptions play a central role in that debate too, with a main subject of contention being precisely whether, if determinism is true, people and other bearers of abilities *can* do otherwise than they in fact do. If incompatibilism is correct, then determinism means that things cannot be otherwise than they are, given the actual past; and in conjunction with CAP, that would entail that things that will not actually happen have no chance of happening (and thus that things which will actually happen must happen, so have chance one of happening). Incompatibilism in this traditional debate leads naturally to incompatibilism about chance. Correspondingly, compatibilism in the traditional debate connects with compatibilism about chance. The connections aren't perfect, because much of the free will debate is concerned with conceptions of human ability stronger than the mere truth of a 'CAN' claim, as here. But they are nevertheless suggestive.

This connection might be severed if one was inclined to dispute whether all outcomes have a chance of coming about. Call *universalism* about chance the thesis that every physically possible event has a chance.[17] Some participants in the free will debate might wish to dispute universalism for reasons to do with human agency; other philosophers may worry about it because they

have a particular view about which features of a setup give rise to chances.[18] But most philosophers haven't been moved by such considerations. I accept universalism, and most parties to the debate over chance and determinism do so also. I think arguments for it can be given, but to avoid digressing, I will simply assume universalism. I will also assume that the case we are interested in is not whether determinism is true, but whether, if true, it must trivialise chances. So from now on I'll systematically neglect the 'libertarian' position in the debate over chance, i.e. the version of incompatibilism holding that there are non-trivial chances and that, therefore, determinism is false.

If one were so inclined, given determinism and CAP, one could exploit the connection between chance and 'CAN' claims in the debate over free will. For example, if one was entirely convinced of incompatibilism about chance, one could argue than any outcome with no chance cannot be produced by anything. Then determinism would undermine the existence of any sort of ability to do otherwise. I'm not so inclined, partly because I'm not convinced of the truth of incompatibilism about chance, but also because I suspect the general debate over the compatibility of determinism and the ability to do otherwise subsumes the debate over chance and determinism. So it strikes me as more appropriate to use broader debate simply as a model to help clarify the issues concerning chance.

4 Simple Incompatibilist and Compatibilist Theories of '*X* CAN ϕ'

The debate over compatibilism generally is vexed to say the least and I can hardly resolve it here. I will, however, illustrate a couple of ways that the debate might go that impacts on my present concerns. I will assume without argument that 'CAN' is a sentential operator, that is, that the correct form of sentences like '*X* CAN ϕ' is 'CAN(*X* ϕs)'; if you prefer a predicate modifier account of 'CAN' you should be able to rework the details without much difficulty.

Incompatibilists give a semantics for sentences like '*X* CAN ϕ' according to which they mean something along these lines:[19]

(4) It is not now settled that *X* does not ϕ.

According to a now standard conception (Earman, 1986), if a world w is *deterministic*, then, given the laws of w, the entire history of w supervenes on the occurrent state of w at any instant (assuming for simplicity a linear time ordering). As such, the initial conditions of a deterministic world, given the laws, are sufficient to settle the entire history, and therefore at every moment, every fact about that world's history is already settled. (I assume that if x supervenes on y, then if y is settled, then so is x.) As such, the only things that can happen, if the incompatibilist analysis (4) is right, are those which will happen. And, if CAP is correct, the only outcomes with

non-zero chance are those which have chance one. This obviously trivialises chance in deterministic worlds. While (4) clearly has some intuitive appeal—witness the pull of incompatibilist arguments against deterministic free will—this incompatibilist theory of 'CAN' manifestly fails to do justice to our common beliefs about (1) and about when it is reasonable to attribute an ability.

By way of contrast, consider this sample compatibilist reading of '*X* CAN ϕ':

(5) *X*'s intrinsic state and immediate circumstances are consistent with ϕ-ing.[20]

According to this view, *X* possesses an ability in virtue of the state of a region including *X* and *X*'s immediate surroundings. Some support for (5) as giving the truth conditions for 'CAN' claims comes from our standard practices in attributing abilities. Ordinarily, we take ourselves to have sufficient grounds to assign an ability when we know the local circumstances and intrinsic state of the proposed bearer. Moreover, if the existence of constraints from outside that region occurs to us at all, we frequently take them to be impediments to the exercise of the ability, rather than showing that no ability exists (we might take the behaviour to be coerced but not deny that it could behave otherwise). If there are no obvious signs of impairment, constraint, or coercion, we are standardly content to regard impeding factors from further afield as irrelevant to whether the ability exists. If an account of the truth conditions of 'CAN' must fit our use, a compatibilist semantics for 'CAN' like (5) fits our naive compatibilist practices tolerably well.

Determinism as characterised in this section is a thesis about the relationship between *global* instantaneous states—the claim that, given the laws, any state supervenes on any other state. To generate tension between determinism and compatibilism about 'CAN', we need the laws and the state of a small region around *X* to fix the global state. That is, we need the thesis I call *unique extension* to be true: that there is a unique way to extend a given region of an instantaneous state to a global instantaneous state. But the thesis of unique extension is not true of most interesting physical theories (it fails for classical mechanics and for relativity, for example). And without unique extension, we get no problem for determinism from the existence of non-trivial compatibilist abilities. So (setting aside the surprisingly vexed issue of whether the theory is deterministic at all), the failure of unique extension in Newtonian physics shows that if abilities do turn out to be trivialised in a theory, they are not trivial *merely* because the theory is deterministic.[21]

The foregoing discussion shows that a simple kind of compatibilism about chance and determinism, flowing from the combination of (5) and CAP, is consistent. But it is important not to misunderstand my argument here. I

am not claiming that the naive compatibilism captured by (5) is the correct account of 'CAN', and I'll argue in the following section that it is semantically inadequate. If you prefer another compatibilist account of 'CAN', you can substitute that throughout—CAP is supposed to be consistent with a variety of stories about what having an ability precisely consists in. The argument in this section is just that there is a prima facie case to be made, on the basis of a simplistic but not absurd compatibilist account of 'CAN', that not every deterministic theory trivialises chance.

Nor am I claiming that every deterministic theory will feature non-trivial chances. I think it likely, in fact, that some deterministic theories really do restrict chances to extremal values. Assume that (5) correctly characterises 'CAN', and consider the thesis we might dub *regional isolation*, which holds in a physical theory just in case the physics permits the state of a region R to screen off influences outside of R. Arguably, regional isolation does trivialise abilities, on the assumption that abilities are analysed according to simple-minded compatibilism—if the immediate circumstances suffice to determine the forward light cone, then there will only be one outcome compatible with those circumstances.

Regional isolation is false of classical physics, because with no upper bound on the velocity of causal signals, no region can be guaranteed to be isolated in the required way (without imposing strong boundary conditions). But special relativity entails regional isolation, because in that theory any Cauchy surface across a light cone determines its forward and backward domain of dependence, so that influences from outside that surface are irrelevant to what occurs in that domain of dependence. This shows that special relativity trivialises chance even under the naive compatibilism expressed by (5). The first conclusion, of course, is that this is not particularly decisive—the simple-minded compatibilism of this section is hardly plausible. And it does turn out that the more sophisticated theory I develop in the following section does permit non-trivial abilities even under the assumption of regional isolation.[22] The second conclusion is that we've now diagnosed what trivialises chances in special relativity, and it is not the thesis of determinism by itself, but rather the related but distinct thesis of regional isolation. Far from undermining the compatibility of chance and determinism, this observation supports it.

5. A More Sophisticated Account of 'X CAN ϕ'

Neither of the two simple theories of 'CAN' just considered are particularly attractive, primarily because neither is flexible enough to do justice to the diversity of behaviour that 'CAN' (and the related 'MUST') exhibit. Specifically, the incompatibilist theory doesn't explain why competent speakers of English should be so frequently willing to judge 'CAN' claims as true; while the compatibilist theory cannot explain the clear intuitive appeal of incompatibilism.

Moreover, the simple compatibilism (5) above faces the problem that not all theories make an objective distinction between immediate circumstances and other circumstances; and even those theories which do, sometimes draw that distinction in the wrong place for the compatibilist.

Thankfully we do not need to adopt either of these naive accounts of 'CAN'. The standard view amongst linguists working in formal semantics is a much more sophisticated and attractive proposal, the idea that 'CAN' is a *relative modality* (Kratzer, 1977; 1981; Lewis, 1979; Portner, 2009). Roughly, this theory (which I'll call the Kratzer-Lewis account) proposes that 'CAN', like 'POSSIBLY', is effectively an existential quantifier over some collection of possibilities—'CAN ϕ' is true iff there is a possibility at which ϕ is true. But, as I mentioned earlier (§1), 'CAN' differs from 'POSSIBLY'. For one thing, there are many uses of 'CAN' which cannot be captured by understanding it as a quantifier over a single domain of possibilities. The Kratzer-Lewis account captures these various senses by varying the *modal base*, the set of worlds that the modal quantifies over. In the case of epistemic modals, the modal base is the set of epistemic possibilities. In the case of interest to us, the dynamic modals, the modal base is a set of physically or metaphysically possible worlds. So one way in which 'CAN' is a relative modality is in being relative to a domain of possibilities.

But merely varying the modal base in this crude way (distinguishing only broad species of possibility) doesn't enable us to capture all the linguistic phenomena associated with 'CAN'. Consider this classic example:

(6) I can speak Finnish. (Lewis, 1976: 77)

Dogs cannot speak Finnish, because dogs don't have the laryngeal structure (nor the brain capacity) to speak any human language. But the facts about my larynx, unlike a dog's, are consistent with my speaking Finnish, so (6) is true. On the other hand, I'd make a terrible translator, since I can't speak Finnish. Contradiction? Not exactly: while my speaking Finnish is compatible with facts about my physiology (i.e. there are objective possibilities where I speak Finnish and those physiological facts remain fixed), it is not compatible with facts about my schooling (so there are no objective possibilities where I speak Finnish without being trained to). When we leave open which facts are being considered, as I just did, we may equivocate over the truth of (6). The possibility of such equivocation is good evidence that an expression is *context sensitive*—that the expression has an implicit parameter, the value of which is fixed by the conversational context and not explicitly by the other parts of the expression.

Whatever context is, it looks like there are many examples in which something which plays the role of context ('conversational score', as Lewis 1979 has it) is required. In this paper I don't intend to take a position on what context is, or say much about how it evolves. I will assume that, in many cases,

explicit use of a proposition p by a conversational participant suffices to ensure that p is from then on presupposed, and that further utterances must be consistent with p (Stalnaker, 1974). This suggests at least that the context at a time will include a *context set* of propositions (pre)supposed (Stalnaker 1998). A lot of the parameters involved in Kratzer's semantics will be fixed by this set of background propositions, but I assume that context has whatever additional features (if any) that are required to fix the remainder of the context-sensitive parameters in order to replicate the observed behaviour of 'CAN' claims.

The potential equivocation over (6) is prevented if we explicitly introduce an 'IN VIEW OF' clause, stating the facts we're assuming in the backgrounds:

(7) a. I can speak Finnish (in view of how my larynx is).
b. I can't speak Finnish (in view of what my schooling involved).

The sentences in (7) are both true, and remove the possibility of equivocation, if we take them as potential completions of (6). These parenthetical glosses may be understood as making explicit mention of the additional finer structure in the modal bases relevant to utterances of (6) in different contexts, as Kratzer (1981) holds, where context additionally provides an ordering of the worlds in the modal base. This ordering suffices to delimit which possibilities count as close enough to be considered as bearing on the truth of the claim. The sentences in (7) are less context sensitive because the features upon which the quantificational restriction depends are explicitly fixed (and context cannot fix them again).[23]

Whatever the relationship between (6) and (7), it is clear that (6) uttered in different contexts can express different propositions, among which are those naturally expressed in a less context-sensitive way by the sentences in (7). On this kind of sophisticated account of the semantics of 'CAN', a sentence of the form 'X CAN ϕ' will, in a given context, express a proposition very close to that expressed by the more contextually stable sentence 'X CAN ϕ IN VIEW OF R', which in turn is true just in case there exists a possible world where X ϕs while under the conditions in the restrictor R—that is, that X's ϕ-ing is consistent with the restrictor conditions.[24]

I wish to emphasize that even though 'CAN' is a relative modality, this does *not* mean that it is in any way epistemic or 'subjective'. In contrast to the epistemic possibilities, relativised dynamic modals are still perfectly objective—there is an objective fact of the matter concerning whether a certain contextual restriction is, or is not, in place with respect to a given claim 'CAN ϕ'; and an objective fact concerning whether that restriction is, or is not, compatible with the proposition ϕ. So the mere fact that 'CAN' is relativised doesn't make it—or the chances that I've argued flow from true 'CAN' claims in virtue of CAP—less than objective (see also p. 24).

6. The Possibility of Deterministic Chance, and the Significance of this Discussion

Given this sophisticated account of 'CAN', CAP entails that 'HAS SOME CHANCE' is a relative modality. Happily, this conclusion dovetails nicely with a burgeoning consensus amongst philosophers that probability is a fundamentally *conditional* notion (Hájek, 2003). Events have different chances relative to different background conditions. If true, this resolves, or at least puts in a new light, several outstanding difficulties with unconditional probability, notably the problem of probability-zero events in the ratio definition of conditional probability, and the reference class problem. One might, moreover, appeal to differing modal bases to explain the difference between different kinds of probability. So the Kratzer-Lewis theory, when applied to chance, has several nice and very plausible consequences. I hasten to emphasize, however, that this convergence with the theory that probability is fundamentally conditional does not amount to equivalence. The chances in the present paper are unconditional chances which turn out to be context-sensitive. The obvious way to account for these phenomena in the fundamental conditional chance view is that all chance is conditional, with the conditioning events set by context when not stated explicitly. But one might alternatively take the present theory as a way of taking the insights of the fundamental conditional chance view and representing them in a framework which does continue to regard unconditional chance as basic. Either way, the present view adds to the consensus that relativising chance to background conditions, however that is implemented, is the best way to respond to a number of outstanding issues.

More germane to our current discussion, however, the Kratzer-Lewis theory in conjunction with CAP holds out the prospect of deterministic chance, since there can be true ascriptions of non-trivial chance even in worlds where determinism is true. One quick way of seeing this is to note that both of our naive accounts of 'CAN', (4) and (5), are potentially expressed by 'X CAN ϕ', with the appropriate restrictors in place. Since there is an admissible compatibilist reading of 'can', it follows that there can be non-trivially true 'CAN' claims and (given CAP) non-trivial chances, even under determinism.

We may actually go further. The naive compatibilist account (5) relied on an imprecise notion of 'immediate circumstances', which was problematic in cases where no physically privileged precisification of that notion existed. The Kratzer-Lewis theory allows for any circumstances whatever to function as a restrictor on the modal, so there is no demand for a physically privileged division.[25] In fact, in light of the apparent vagueness of 'immediate circumstances', a contextualist account of 'CAN' is arguably the only one capable of explaining our ordinary use of that expression. So, unlike the naive theory, contextualist compatibilism does not rely on a physically privileged notion of 'immediate circumstances'.[26] As I argued earlier, (5) gets some support from

the fact that we are normally perfectly content to attribute an ability on the basis of an absence of obvious internal incapacity or external constraint (see also p. 10). No purely incompatibilist account of 'CAN' could explain this without invoking a widespread error theory. But what is relevant to an ability attribution seems to vary from context to context, and the naive theory (5) cannot explain this. The Kratzer-Lewis theory, even without an explicit account of what makes some condition relevant in a context, can explain these successes of the naive theories without the difficulties.

Time now to draw my conclusion most explicitly. Consider this sentence:

> **IC** If a world is deterministic, then no possible outcome in that world has any chance there other than 1 or 0.

IC is not about language, and incompatibilists who use IC to express their view are not intending to make any claim about language. But if CAP is right, the foregoing linguistic considerations do have a bearing on evaluating it. For if those considerations are correct, the sentence IC has a subtle context-sensitivity, due to its containing the phrase 'HAS ANY CHANCE'. Thus, which proposition is the meaning of IC will vary from context to context. Given that IC is the sentence incompatibilists would use to present their position, compatibilism will be defensible if there exists some context in which the proposition expressed by IC is false. In such a context it would be entirely appropriate to say that determinism doesn't trivialise ability, and does not trivialise chance. In such contexts, any argument which relied on a claim like IC could be challenged as unsound, even if that sentence expresses a truth in other contexts.

I think it is quite clear that such contexts exist, and that therefore it would certainly be premature to reject compatibilism on the basis that there exist other contexts where IC is true. My own view is that compatibilism is the correct attitude in ordinary contexts, largely because ordinary unreflective use is clearly compatibilist, and there is no evidence that uses of 'CAN', and thus 'HAS SOME CHANCE', invariably force the whole past and the laws to be relevant in a context.[27] The hardline incompatibilist, remember, has a severe task in front of them: to show that, even though 'HAS SOME CHANCE' is a relative modality, it is *never* the case that the contextually salient facts permit more than one chance outcome. Moreover, this must be the case because of determinism (and not some other source). So it looks like the incompatibilist must argue that the facts of determinism must always be contextually relevant, as must sufficient facts about some past state of the universe to trivialise future ability claims. This is what it would take for IC to always express a truth. This just seems highly implausible.

Yet if what I've argued is correct, it would be equally unreasonable to expect that IC expresses a falsehood in every context. In fact, in contexts where the past and laws are salient and thus contextually relevant, we should expect

IC to express a truth. And this is exactly what we see, as predicted by the Kratzer-Lewis theory. IC is spontaneously professed largely in philosophical contexts, precisely where issues of determinism have been introduced and discussed.[28] In such contexts, with determinism in the conversational background, the facts to which the modality is relativised will plausibly expand to include all facts about what is presently settled (in the sense of (4)). In these cases, the incompatibilist reading of 'HAS SOME CHANCE' is dominant; and this explains why incompatibilist intuitions can be so readily elicited in discussions of determinism. So while 'THE COIN HAS A 50% CHANCE OF LANDING HEADS IF TOSSED' can be true even in deterministic worlds, 'THE COIN HAS A 50% CHANCE OF LANDING HEADS IF TOSSED IN A DETERMINISTIC WORLD' cannot be true, or at least not in ordinary contexts. But outside of such contexts, it is plausible to suppose that the background facts to which 'HAS SOME CHANCE' is relativised include just immediately salient facts about the object under discussion and its surroundings. In such contexts, a compatibilist reading of 'HAS SOME CHANCE' is readily available and that explains why ordinary speakers are, most of the time, content to assent to many chance claims that the incompatibilist finds problematic.

*

So much for my sympathies towards compatibilism, especially in the Kratzer-Lewis form. If you share my sympathies, there is ample reason here to suspect that the philosophical near consensus is premature: compatibilism about chance is a viable and attractive position concerning chance ascriptions in ordinary contexts. But, perhaps counterintuitively, the key question for both compatibilists and incompatibilists doesn't concern the truth of IC; in the present framework I think it is fairly easy to generate ordinary contexts where it is false, and philosophical contexts where it is true. The really important point, and the *main point* I'd like readers to take away from my discussion—a point which is equally as important for those with incompatibilist sympathies—is that the truth or falsity of IC in ordinary contexts is not nearly as important as whether IC manages to play the role required of it in philosophical arguments.

For, like most philosophical claims, the interest and importance of IC comes from the arguments in which IC features as a premise. The present discussion is crucial here, because the hidden context sensitivity I've argued exists in IC and other chance ascriptions could mean that in some argument about chance, even if all the premises are defensible and the argument appears valid, there may be no single context in which the premises are true along with the conclusion. I will consider one important example, expressed in this quote from Loewer:

> [W]e have a paradox—*the paradox of deterministic probabilities*. If, as Popper and Lewis claim, objective probabilities cannot co-exist with deterministic dynamical

> laws then the probabilities that occur in [classical statistical mechanics] are subjective. But if that is so then, as we have just argued, these probabilities do not ground the lawfulness of principles that are usually taken to be lawful. Loewer (2001: 612)

An argument is here offered to a paradoxical conclusion, in which IC, or something near enough, is a crucial premise. The other premise is that certain robust generalisations—such as that ice cubes placed in warm water melt—are genuinely laws (supporting counterfactuals and being confirmed, typically, by their instances). Yet, and this is the paradoxical conclusion, it seems they cannot be laws without genuine objective chances (at least while we offer the normal explanations in classical statistical mechanics as to why these generalisations obtain (Albert, 2000)).

Is Loewer's argument genuinely paradoxical? I think not (for the record, neither does he—but for different reasons). For in contexts where the first premise, IC, expresses a truth, I think it is difficult to see how the second premise can be true. Such contexts, I just argued, involve the contextual salience of the fundamental deterministic laws. As Loewer admits, generalisations like 'ice melts in warm water' are not entailed by the laws. It is difficult, at best, to maintain that some generalisation L is a law while simultaneously noticing that $\neg L$ is consistent with the highly salient fundamental laws. This, I suggest, makes the second premise false in such contexts. But in contexts where the second premise is true, the fundamental laws can't all be salient; and without the salience of these laws, it is difficult to get IC to express a truth. Even though both premises of the paradox of deterministic probabilities can be true, and are individually defensible, they are not jointly defensible or jointly assertible. I think this suffices to resolve the paradox, and to show the merits of the present proposal about chance and its relation to determinism. These merits lie in bringing to our attention subtle features of the arguments we offer about the metaphysically important notions of chance and determinism. Attention to language doesn't resolve the truth of IC; but it does ward off fallacious inferences we might make in the process of discussing it, in particular, the fallacies of equivocation that arise from ignoring the role of context. (See also Williamson's remarks quoted in Appendix A.)

7. Further Objections to CAP

The foregoing discussion, connecting chance and ability ascriptions, does emphasize three continuing difficulties with CAP.

Firstly, one might object that connecting chance with abilities does more to highlight the differences between chance events and the outcomes of free actions than their similarities. We can sharpen this into a couple of potential counterexamples to the proposal:[29]

Left-to-Right Consider Pokey, the world's worst darts player, who has no control of the dart whatsoever: he just throws at random. Because of that, it seems he has some chance of hitting the dartboard—it would hardly be random if he threw darts anywhere but there. Yet there is some pressure to deny that he can hit it, because without the exercise of control we are reluctant to ascribe ability.

To this case, and others like it, the response that appeals to me is to say that while control might be important for some senses of ability, it isn't important for all senses. There is apparently what we might call the 'try-conditional' sense of 'CAN', according to which 'X CAN ϕ' means something like (Thomason unpublished):

(8) If X tried to ϕ, X would succeed.

Clearly *this* sense of 'CAN' cannot be used when X is not in control of his actions; it is certainly not the case that Pokey would succeed in hitting the dartboard if he tried to. But while this is an important sense of 'CAN', it is not the only sense.[30] In particular, this is *not* the sense of 'CAN' which features in CAP—to see this most clearly, note that if 'CAN' is read in this sense, the sentences in (1) aren't false—or even troubling! For since the counterfactual 'IF WE HAD TRIED TO TOSS THE COIN TO LAND HEADS, IT WOULD HAVE LANDED HEADS' is false, the try-conditional analysis of 'CAN' in (8) entails that the first conjunct of (1a) is false, and not in conflict with the second conjunct at all. Yet this prediction is absurd: for it to be true that a coin can land heads, it need not be the case that if we flipped it, it would land heads. So to explain the badness of (1), which is the primary rationale behind CAP, we require a weaker sense of 'CAN' than the try-conditional sense, which the 'compossibility' sense given by the Kratzer-Lewis semantics is well placed to provide.[31]

Right-to-Left I argued earlier that (6) can be true. CAP then entails that I have some chance of speaking Finnish. But it seems downright strange to say that, even though I don't know how to speak Finnish, I have some chance of speaking Finnish.

I think the best response to this is attack: if I can speak Finnish, then I do have some chance of speaking Finnish. For think about how I might actualise my potential to speak Finnish: I go to Finnish night school, I study hard, I move to Helsinki, etc. All of these courses of action are open to me; and there are physical processes I could initiate which result in them being actual. I am, of course, unlikely to carry them through: but there is some non-zero chance that I do initiate these actions. Insofar as there is a possible future for me in which I do speak Finnish, that possible future has some chance of being actual (making use here of the BCP). Of course even if I can speak Finnish in this sense, it doesn't mean I can *right now*; and if one reads 'THERE IS SOME CHANCE OF MY SPEAKING FINNISH' as entailing that I might all of a sudden start speaking Finnish, it is false. But that reading is hardly the only one available.

The objector might retrench at this point: the compatibility of my speaking Finnish with the state of my larynx and language module could, it seems, remain even if I was certain *not* to carry through any course of action involving

> learning Finnish, and remain even if I will die in the next minute. Then there is apparently an ability but, given my impending death, no chance. Once again, I'm not so sure. When it really is salient that one is about to die, I don't share the intuition that one can speak Finnish: one can't now, and won't have the chance to learn. Or, if we do persist in having the intuition, we must regard it as compatible with how my larynx is now that I speak Finnish now, and hold that there is some (miniscule though it may be) chance that the particles align to allow me to utter, and utter with comprehension, a sentence of Finnish. Put it another way: if I really have no chance of speaking Finnish because I've no time left to do so, this must be something included in the background context (since it is not in general true that 'I WILL NOT ϕ' and 'I HAVE SOME CHANCE OF ϕ-ING' are incompatible). But, as I've argued, the relevant 'CAN' claim for comparison is one that involves the very same contextual background. And if it is assumed that I will do nothing more because I will die, I fail to see how my speaking Finnish is compatible with my not doing anything further (and in particular not speaking Finnish). On the other hand, if what grounds the 'CAN' claim is some other background, I see no reason to think that there is no chance relative to that other background. Once again, careful attention to context, in order to avoid equivocation, dissolves this apparent problem.

Secondly, a problem emerges from the persistent temptation to think that, even if 'CAN' is a relative modality, some relativisations are better than others. This temptation exists because the expansion of relevant facts engendered by discussion of determinism (from including purely local facts to including global facts and laws) is not readily reversed. Lewis explains this in terms of his 'rule of accommodation'—it is fairly easy to induce a conversational shift to a richer set of background facts but hard to induce a shift back to a more impoverished set of facts. It may well be that consciously ignoring formerly salient facts is psychologically more difficult than remaining oblivious to not-yet-salient facts. Be that as it may, must we think that, in Lewis' words (1979: 247), 'what is true with respect to the outward-shifted boundary must be somehow more true than what is true with respect to the original boundary'? In a neat answer to this question, Lewis continues:

> I see no reason to respect this impression. Let us hope, by all means, that the advance toward truth is irreversible. That is no reason to think that just any change that resists reversal is an advance toward truth.

Thirdly, I wish to return briefly to the putative platitudes discussed in §2. At least some of the arguments there turned on ordinary beliefs about ability. There is no real reason to expect any of those ordinary beliefs to be overturned once we see the correct semantics of ability ascriptions; after all, the evidence for the semantic theory is in part just those ordinary judgements about 'CAN' claims. But one may be worried that if 'CAN' claims and 'CHANCE' claims are context sensitive, then some of those principles may not turn out to be so obvious after all. I've already flagged a possible concern about the

Future Principle, so let me turn to that first. The worry is that, once we admit cases of backward time travel, it will in some contexts turn out that the past can be changed, and that it therefore has some chance of being different; which might be thought to conflict with the platitude that what's past is no longer chancy. However it is not clear to me that this conflict is genuine. It will be true in almost all contexts that the past can be held fixed, and what varies is how much of the past is contextually salient. The debate here is over whether there are any contexts in which the actual outcome of a process may be properly neglected. I myself am sympathetic to the idea that in some special and extremely rare contexts we can treat such outcomes as if they are not fixed, and ascribe non-trivial chances to past outcomes and posit non-trivial abilities to produce those outcomes. The rarity of such contexts it seems to me is enough to explain the intuitive appeal of the future principle. Moreover the existence of exceptions to the future principle in cases of backwards time travel seems to me right; I think it is fair to say that a coin tossed by the time traveller as she enters the time machine has a chance of landing heads 50 years earlier, because in such contexts the right facts to hold fixed are those in the past of the time traveller's personal time. If her tossing the coin now can cause its landing heads 50 years earlier, I see no obstacle to saying that it also gives rise to the chance then.

A related concern may be raised about the Principal Principle: how can we manage our belief in accordance with the PP when in different contexts different chances will obtain? Surely what we believe doesn't change from context to context? The brief answer here is: there is no conflict. So long as credence is a propositional attitude, we will believe the same propositions in every context. What will vary are the sentences used to express our beliefs, and it may well be that we can use a chance ascription in one context to express a given belief but be unable to use that sentence in another context to express that belief. This may give the appearance that we change our mind when the context changes, but is in fact no more a case of changing one's mind than when one uses 'Today is Wednesday' to express your belief today, and 'Yesterday was Wednesday' to express that very same belief tomorrow. If there are other apparent difficulties produced by the interaction of context sensitivity and others of the platitudes, I am confident they can be similarly dissolved.

Conclusion

My main goal here was modest. By indicating how, in accordance with plausible principles about chance, determinism isn't necessarily trivialising, I aimed to show how attention to the semantic details can illuminate and clarify the apparently metaphysical thesis of incompatibilism and its consequences—when that thesis is actually used in particular arguments in particular contexts. My secondary goal was to introduce readers to the ability constraint on chance, point out its initial plausibility, and respond to some basic

objections. Whether CAP as sketched here can be more fully articulated and defended, I intend to consider elsewhere.

Acknowledgments

Research on this paper was carried out with the aid of research leave funding from the Arts and Humanities Research Council (under grant AH/E003184/1). The paper benefitted from discussion in a philosophy of probability reading group at Oxford and from a presentation to the philosophy department at the University of Melbourne. Considerably earlier versions of the paper were presented at the 2006 meeting of the AAP in Canberra; Oxford; the University of Illinois Chicago; and at the 2005 Formal Epistemology Workshop in Austin. Thanks to members of those audiences, and to John Burgess, John Cusbert, Josh Dever, Adam Elga, Bas van Fraassen, Luke Glynn, Alan Hájek, Hans Halvorson, Toby Handfield, Aidan Lyon, Maria Lasonen-Aarnio, Anna Mahtani, Jeff Speaks, Barbara Vetter, Tim Williamson, and Alastair Wilson, for discussion and comments. I owe particular thanks to Jonathan Schaffer and an anonymous reviewer for this journal for their very useful comments on a previous draft.

Appendix A: Must Chances Be Probabilities?

I argued in §1 that CAP is plausibly compatible with the standard probability calculus. The official line in this paper is that chances which obey CAP are probabilities. But is this a requirement on chance?

The standard Kolmogorov axiomatisation is an attempt to capture the concept of probability, but cannot plausibly be claimed as a stipulative definition: for one thing, there were people making use of the concept of probability long before the formal axiomatisation came along. A formal system is good to the extent that it captures in a precise way what is true of the informal concept. It is not in itself a problem that some proposed theory of probability violates the formalism: what is objectionable is a theory which violates the concept that the formal theory was intended to capture. This is clearest in cases where a theory of probability violates the formalism without doing injustice to the underlying concept, and in such cases alterations to the formalism have been proposed and drawn widespread approval. Consider the recent example of those who maintain, contrary to the Kolmogorov theory, that some conditional probabilities $\Pr(A \mid B)$ are well-defined even when $\Pr(B) = 0$ (Hájek, 2003). The evidence for their claim is that many people agree on the values of, and do not baulk at assigning values to, certain of these conditional probabilities, and this evidence drawn from common belief about probabilities is sufficient to make us hesitate to regard the standard formalism as perfectly capturing the concept.

In §1, I provided evidence of a similar inadequacy in the Kolmogorov formalisation of the concept of chance. The evidence is just that competent

speakers, deploying common belief about chances, believe that sentences like those in (1) are inconsistent. Since some sentences like (1a) can be true on the standard formalism, we have a mismatch between our judgements about chance and the predictions of the formal theory. As in any such case of a mismatch we need to give an argument as to which we should abandon—it is not automatic that we abandon our common belief in favour of the formalism, as the error could easily have been introduced exactly when that formal system was formulated. This is not to say that the Kolmogorov axioms are not useful or approximately correct; the present theory entails that in cases other than chance zero events, the axiomatisation gets it right.[32] It *is* to say that we cannot take the Kolmogorov theory as a stipulative definition.[33]

Nor can we simply reject common belief. Too many philosophers would at this point maintain that we should reject the old concept of chance in favour of the new precise scientific concept correctly described by the Kolmogorov axioms. Of course we could; but the fact that the older concept is less simple and manageable than the readily axiomatised precise concept is hardly evidence that chance is incoherent, or doesn't apply to anything, or has no useful theoretical role. I'm not comfortable with the kind of massive error theory about ordinary English uses of the familiar word 'CHANCE' that this revisionary attitude would result in. I'm much more inclined to trust the judgements of reflective competent English speakers on (1) than revisionary philosophers in the grip of some received formalism.

Of course the prospect of convicting the folk of widespread error will not trouble the revisionary philosopher. The real problem arises when the revisionary philosopher tries to convince the folk of some philosophically interesting consequences of their view. After all, the word 'CHANCE' had a life outside of mathematics well before Kolmogorov came along, and continues to feature in arguments of philosophical interest now. If 'CHANCE' isn't synonymous with 'OBJECTIVE KOLMOGOROV PROBABILITY', these arguments involving 'CHANCE' aren't going to be correctly captured by formalising them as if they are synonyms! We must be sure that in revising our use, we don't revise away all the philosophical interest in the term; and the only way to be sure of this is to pay close attention to how the term actually functions.

Examples are invidious, but this kind of problem pervades some of the more radical conclusions drawn from contemporary philosophy of physics; in its apparent eagerness to revise common opinion to fit with current science, there is a very real danger of simply mistaking the subject matter of the relevant arguments. I once heard a prominent physicist maintain that, since particles do not have precise positions according to quantum mechanics, that no object is located anywhere. This rests on a very naive view of how the position quantity of quantum mechanics maps on to the ordinary concept of location, and only insufficient attention to and respect for the ordinary concept could have led someone into making this implausible claim. Tim Williamson makes the similar point that those

> who most disdain language are the most likely to be its victims. Again, those who neglect logic in order to deduce philosophical results from natural science make frequent logical errors in their deductions; their philosophical conclusions are not logical consequences of their scientific premises. For example, some supposed tensions between folk theory and contemporary science depend on fallacies committed in the attempt to draw out the consequences of common sense beliefs. (Williamson 2008, 46)

I think another, similar, error would be to reject an otherwise attractive proposal about 'CHANCE' on the basis that it violates the probability calculus. That it might not fit with a simple and attractive story about probability may count against a view, but is not on its own conclusive. For the fact is that the Kolmogorov analysis fails to do justice to simple and straightforward common beliefs about chance, captured in (1). This is similarly evidence for my view and against the adequacy of the Kolmogorov axioms as an account of chance. And this suffices to call (3) into question, and provide another way of blocking the obvious objection to CAP—for someone inclined to take it.

Notes

[1] False, that is, unless the situation is extraordinary, perhaps in virtue of someone securing the coin to the table after the utterance of the first part of (1a) or in virtue of a change to the conversational context of utterance mid-stream. As we will see below (§5), I will suggest that both 'CHANCE' and 'CAN' sentences are context-sensitive.

[2] 'Dynamic', from Greek *dunamis*, 'power', and hence ability. The contrast is with epistemic and deontic senses of 'CAN', as in 'YOU CAN'T BE SERIOUS!', and 'YOU CAN ALL GO HOME NOW' respectively.

[3] Indeed, as propensity theories essentially claim to account for chance by means of dispositions, it might be argued that all of them give a central role to this connection.

[4] As such chances plausibly retain a connection with dispositions, as dispositions are closely connected with habituals (Fara, 2005).

[5] It is also worth noting, I think, the crucial and distinctive role played by the noun phrase $\ulcorner X \urcorner$ even in CAP, denoting what others have called the 'chance set up' Hacking (1965). In emphasising this, even this simple version of CAP demonstrates its potential: firstly, it suggests a partial resolution to the notorious *reference class problem* for chances (Hájek, 2007), by making chances depend on the reference class determined by specification of suitable trials of the chance set up. Secondly, it illustrates the rather impoverished conception of chance that Schaffer (2007: 115) uses, following Lewis (1980: postscript C)—it may be that specifying a world, proposition and time do not suffice to yield a unique chance function without further specification of the chance set up.

[6] And though it connects chance and (one kind of) possibility, it does not entail anything like a principle of indifference according to which any outcome which can happen should be assigned the same chance as any other.

[7] Claiming, as I do, that CAP is a truth about chance is far from claiming that there exists a reductive analysis of chance in terms of ability. CAP is a truth about chance that should feature in any good theory of chance, regardless of whether that theory takes the form of a conceptual analysis. I would note however that an outcome is highly probable to the extent that it does happen when it can, and it does seem that chance has some connection with how *typical* it is for X to ϕ (one further reason why attention to the habitual quantification appearing (1) is needed). This in turn is surely connected to frequencies of outcomes in cases which share the

same basis for the ability as the actual case. But nothing I say in this paper depends on the prospects for an 'ABILITY THEORY' of chance; CAP remains a truth about chance regardless.

[8] See also the discussion at footnote 25.

[9] Mellor, in defending his account of the connection between chance and possibility, suggests something similar: 'Particulars cannot have point values of continuous quantities over which they have chance distributions' (Mellor, 2000: 26)—though I suspect that by 'CANNOT' he means that it is impossible.

[10] This last observation defuses the potential worry that CAP is contingent, since in every world, every event which can occur in that world has non-zero chance, even if some of those events have zero chance according to the the best theories of chance of other quite distant worlds.

[11] One could respond further by pointing to the attractiveness for actual physics of certain views of spacetime which don't involve points at all (Arntzenius, 2003: 2008); I'm not inclined to rest much weight on this since I'm concerned about the existence of ordinary putative counterexamples to CAP, and resolving them by means of unusual physics misses the point of the objections.

[12] As a reviewer pointed out, things are slightly trickier here than they may seem. Bearing in mind the argument I propose below (§5) that ability attributions are context sensitive, it may well be that in his philosophical context, with the possibility of determinism very salient, Schaffer's requirement that '*p* is compossible with the circumstances' may indeed express a claim that is only satisfiable under indeterminism. My objection is not that there are no contexts where the intuition behind the Realization Principle is correctly articulated formally as Schaffer articulates it. The objection is rather that this articulation is not correct in every context, and particularly that it is not correct in ordinary non-philosophical contexts, where an outcome might be consistent with the relevant circumstances even if not consistent with the laws and total history. More on this below.

[13] For reference, the formal version is 'Suppose $x > 0$ and $Ch_{tw}(A) = x$. Then A is true in at least one of those worlds w' that matches w up to time t and for which $Ch_t(A) = x$.' (Bigelow et al., 1993: 459)

[14] Things are actually rather more complicated than this once we consider cases of time travel. Under the semantics of 'CAN' I eventually endorse in §5, it is in some contexts true to say that the past can be altered, and given CAP it will be equally true in those contexts that the past has some chance of being different (Lewis, 1976). Cases like this include the sense in which one can kill one's grandfather before he reproduces, since one is a good shot with ample opportunity and desire; if this is so, then there is some chance that one will kill one's grandfather. These abilities and the concomitant chances depend on it being appropriate in context to set aside some relevant knowledge, namely, that the grandfather is such that one will fail to kill him. More on this issue below, page 24.

[15] If there are indeterministic theories without explicit chances—as Earman (1986) and Norton (2007: 24) have argued with respect to classical mechanics—then these future possibilities aren't the sort that give rise to true 'CAN' claims, as discussed above (§1).

[16] What the metaphysics of chance and ability are like if 'CHANCE' and 'ABILITY' turn out context-sensitive is an interesting issue that I don't have space to address here.

[17] Perhaps even 'that chance of truth applies to any proposition whatever' (Lewis, 1980: 91).

[18] So von Mises, (1957: 10–2) rejects universalism because chance only correctly applies to 'mass phenomena'; Hoefer (2007) seems a modern descendent of von Mises' frequentism, rejecting universalism because some setups don't have stable enough outcome patterns to support well-defined chances.

[19] The incompatibilist must insist on a special reading of 'IT IS NOT NOW SETTLED THAT *p*', something like: the laws and history up to the present do not entail either p or $\neg p$. I suppose that some such reading can be given, since I think it is plausible that (at least in the philosophy

room) the incompatibilist should be able to express the proposition about the semantics of 'CAN' they endorse. I will here adopt the convention that the form of words in (4) expresses this incompatibilist view, at least in the present (philosophical) context. (But I am not simply yielding (4) to incompatibilists; if compatibilism is right, it seems there will be uses of 'IT IS NOT NOW SETTLED THAT' which express something compatibilistically acceptable.)

[20] Some have objected even to this kind of compatibilism, noting that if desires are intrinsic, then even though an agent could do otherwise than she desired, still, her desires aren't consistent with that other course of action. So, for example, since I don't desire to eat ice-cream, it is not consistent with my desires that I do; and yet it seems I still can. In reply, one may begin by questioning whether desires are really intrinsic, as their content is arguably not fixed by my intrinsic state. But the main response is that even if desires are intrinsic, the existence of *akrasia* seems to suggest that desiring to ϕ is not genuinely inconsistent with not ϕ-ing. Yet as will become clear in the next section, I won't defend (5) except to the extent that there are some contexts where the kinds of facts it claims abilities invariantly depend on, are the facts held fixed.

[21] Of course determinism will place severe constraints on the *exercise* of abilities attributed by this compatibilist theory of 'CAN'; but that the abilities are there, however masked, is what the compatibilist is interested in.

[22] The more sophisticated account of 'CAN' described in section 6 permits abilities to be relativised to arbitrary sets of conditions, not just those which are truths about some local (or otherwise physically privileged) region. So even if, in view of the properties of R and the truth of regional isolation, there is only one possible outcome, that doesn't mean that there is only one possible outcome in view of any set of conditions whatever. The case then needs to be made that these other restrictors are in play when we use 'CAN'; this is not as difficult as it might seem, as the foreignness of regional isolation to most native speakers shows that it is not a contextually salient condition in their uses of 'CAN'.

[23] Alternatively, the explicit glosses in (7) might be understood as simply making overt the covert structure of (6), in which case (6) would really be elliptical, and 'CAN' would be semantically elliptical for 'CAN IN VIEW OF F', where 'F' denotes a set of conversational background facts to which the modality is relativised. In this alternative we need not introduce the ordering structure in addition to the coarse-grained modal base and the conditions imposed by the restrictor 'IN VIEW OF'. The alternative theory is notably simpler, though the ability of the ordering semantics approach to capture other aspects of modality (for example, comparative possibility) means that the first theory is generally preferable (Portner, 2009: 67–73).

[24] Here the conversational background restrictor provides an *ordering source* (Kratzer, 1981: §3) on the worlds in modal base (in this case, the physical possibilities), such that 'CAN ϕ' is true at a world w iff ϕ is true at at least one world in the modal base v which is (or is close to) ideal at w with respect to the background facts in R—i.e. where those facts are true in a normal or typical way. The additional complexities raised by using an ordering source rather than simply restricting the quantifier to those worlds of the modal base where the restrictor holds don't matter for my purposes here, so in the main text I'll use the simpler model. Further discussion of ordering semantics can be found in Portner (2009: 56–85). I also neglect the further refinements in Kratzer (1991).

[25] It even allows the alethic modality 'POSSIBLY' to be treated as the special case where the restrictor is empty. Again, this shows 'POSSIBLY', in its standard philosophical uses, isn't synonymous with 'CAN'.

[26] So I am *not* suggesting that there is deterministic chance of some event only when there is a privileged distinction between 'local' and 'non-local' regions, and where the local region does not entail the (non-)occurrence of the event. Nor I am suggesting that, if there is any region around X such that outside influences have a sufficiently negligible effect, that X has no ability to do otherwise than it does—for the existence of *one* region that would trivialise abilities does not entail that *every* region does so (a point I return to below, p. 24).

[27] Nor must they invariably induce a context in which the whole local region and the truth of regional isolation are relevant.

[28] Lewis' own declaration in favour of Incompatibilism (Lewis, 1980: postscript B) is a perfect example: it comes after two full pages of discussion of determinism, and this discussion would surely have sufficed to make the supposition of determinism vivid in context.

[29] I owe both of these examples to Jonathan Schaffer, though the first owes a debt to Kenny (1976).

[30] Though I'm not certain of this, perhaps the try-conditional sense can be plausibly accommodated by the Kratzer-Lewis theory. The try-conditional sense of 'CAN' is motivated by considering the use of ability ascriptions in situations involving planning and deliberation. If the mere possibility of success in an action entailed the truth of an ability claim, this would yield implausibly many viable courses of action. The problem seems to be that sometimes we want ability to guarantee success, not just be consistent with it. I think a restrictor-based account can address this issue, making use of the notion of an *easy possibility*. In deliberation, we really only consider actions we may readily and successfully undertake: that is, things we can *easily* do. Those actions then must be drawn from the field of easy possibilities, those possibilities conditional on my action that aren't reliably excluded (Williamson, 2000: §5.3). In urgent cases, like that of an agent deliberating about which tree she can climb to escape a bear, it is clear that the relevant alternative courses of action she should consider are not merely easy possibilities, but those easy possibilities such that their not holding as a result of his action *isn't* an easy possibility. If ϕ is an easy possibility and $\neg\phi$ is not easy, call ϕ a *practically certain* result of action. In urgent deliberation, then, we should consider what we can do in view of what is practically certain for us. In *that* sense, 'Prairie can (in view of what's practically certain for her) climb that tree' is true just when if she were to try, she'd succeed. This is precisely the intuition Thomason rests his account on, but there is no need to go beyond the bounds of Kratzer's theory to accommodate it. Indeed, thinking about things in this way brings out a problem for Thomason's account: if both ϕ and $\neg\phi$ are easy possibilities,then it seems appropriate to rely on 'I can ϕ' in planning action, but certainly not appropriate to think 'IF I TRY TO ϕ, THEN ϕ': for it can easily be that $\neg\phi$. Another problem for Thomason's account is the behaviour of 'CAN' claims under negation, where the predictions of the try-conditional theory seem seriously awry.

[31] Similar remarks apply to the accounts of 'human ability' given by Brown (1988) and Horty (2001), both of which agree with the 'try-conditional' account of 'CAN' in requiring that an ability involves something like an agent's choosing some action, and that action then guaranteeing a certain outcome. Portner, (2009: 202–3) gives a brief argument to the effect that Kratzer's notion of a 'good possibility' behaves similarly to these accounts, and indeed to Thomason's, in quantificational structure, and thus complements the argument of the previous footnote.

[32] In light of the discussion of §6, the convergence between the present theory and those who defend primitive conditional probability is no accident.

[33] That said, the present theory proposes a fairly conservative revision of the probability calculus; other who have claimed that chances are not standard probabilities have often opted for theories that aren't recognisably varieties of probability at all (Fetzer, 1982). Lange, (2006: 399) discusses a chance function that is equal to the corresponding probabilities when defined, but goes undefined in crucial cases.

References

Albert, David Z. (2000), *Time and Chance*. Cambridge, MA: Harvard University Press.

Arntzenius, Frank (2003), 'Is Quantum Mechanics Pointless?' *Philosophy of Science*, vol. 70: pp. 1447–57.

——— (2008), 'Gunk, Topology, and Measure'. *Oxford Studies in Metaphysics*, vol. 4: pp. 225–47.

Bigelow, John, Collins, John and Pargetter, Robert (1993), 'The Big Bad Bug: What are the Humean's Chances?' *British Journal for the Philosophy of Science*, vol. 44: pp. 443–62.
Brown, M. (1988), 'On the Logic of Ability'. *Journal of Philosophical Logic*, vol. 17: pp. 1–26.
Clark, Peter (1987), 'Determinism and Probability in Physics'. *Proceedings of the Aristotelian Society, Supplementary Volume*, vol. 61: pp. 185–210.
Earman, John (1986), *A Primer on Determinism*. Dordrecht: D. Reidel.
Fara, Michael (2005), 'Dispositions and Habituals'. *Noûs*, vol. 39: pp. 43–82.
Fetzer, James H. (1982), 'Probabilistic Explanations'. In *Proceedings and Address of the Philosophy of Science Association*, vol. 2, pp. 194–207.
Glynn, Luke (2010), 'Deterministic Chance'. *British Journal for the Philosophy of Science*, vol. 61: pp. 51–80.
Hacking, Ian (1965), *The Logic of Statistical Inference*. Cambridge: Cambridge University Press.
Hájek, Alan (2003), 'What Conditional Probability Could Not Be'. *Synthese*, vol. 137: pp. 273–323.
——— (2007), 'The Reference Class Problem is Your Problem Too'. *Synthese*, vol. 156: pp. 563–85.
Hoefer, Carl (2007), 'The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance'. *Mind*, vol. 116: pp. 549–96.
Horty, John F. (2001), *Agency and Deontic Logic*. Oxford: Oxford University Press.
Ismael, Jenann (2009), 'Probability in Deterministic Physics'. *Journal of Philosophy*, vol. 106: pp. 89–109.
Kenny, Anthony (1976), 'Human Abilities and Dynamic Modalities'. In J. Manninen and R. Tuomela (eds.), *Essays on Explanation and Understanding: Studies in the Foundations of Humanities and Social Sciences*, Dordrecht: D. Reidel, pp. 209–32.
Kratzer, Angelika (1977), 'What 'must' and 'can' Must and Can Mean'. *Linguistics and Philosophy*', vol. 1: pp. 337–55.
——— (1981), 'The Notional Category of Modality'. In Paul Portner and Barbara H. Partee (eds.), *Formal Semantics: The Essential Readings*, Oxford: Blackwell, pp. 289–323.
——— (1991), 'Modality'. In Arnim von Stechow and Dieter Wunderlich (eds.), *Semantics: An International Hand- book of Contemporary Research*, Berlin: de Gruyter, pp. 639–50.
Krifka, Manfred, Pelletier, Francis Jeffry, Carlson, Gregory N., ter Meulen, Alice, Chierchia, Gennaro and Link, Godehard (1995), 'Genericity: An Introduction'. In Gregory N. Carlson and Francis Jeffry Pelletier, (eds.), *The Generic Book*, Chicago: University of Chicago Press, pp. 1–124.
Lange, Marc (2006), 'Do Chances Receive Equal Treatment Under The Laws? Or: Must Chances Be Probabilities?' *British Journal for the Philosophy of Science*, vol. 57: pp. 383–403.
Levi, Isaac (1990), 'Chance'. *Philosophical Topics*, vol. 18: pp. 117–49.
Lewis, David (1976), 'The Paradoxes of Time Travel'. In *Philosophical Papers*, vol. 2, Oxford: Oxford University Press, 1986, pp. 67–80.
——— (1980), 'A Subjectivist's Guide to Objective Chance'. In *Philosophical Papers*, vol. 2, Oxford: Oxford University Press, 1986, pp. 83–132.
——— (1983), 'Scorekeeping in a Language Game'. In *Philosophical Papers*, vol. 1, Oxford: Oxford University Press, 1983, pp. 233–50.
——— (1994), 'Humean Supervenience Debugged'. In *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, 1999, pp. 224–46.
Loewer, Barry (2001), 'Determinism and Chance'. *Studies in History and Philosophy of Modern Physics*, vol. 32: pp. 609–20.
——— (2004), 'David Lewis' Humean Theory of Objective Chance'. *Philosophy of Science*, vol. 71: pp. 1115–25.

Mellor, D. H. (2000), 'Possibility, Chance and Necessity'. *Australasian Journal of Philosophy*, vol. 78: pp. 16–27.

von Mises, Richard (1957), *Probability, Statistics and Truth*. New York: Dover.

Norton, John D. (2007), 'Causation as Folk Science'. In Huw Price and Richard Corry (eds.), *Causation, Physics, and the Constitution of Reality*, Oxford: Oxford University Press, pp. 11–44.

Palmer, F. R. (2001), *Mood and Modality*. Cambridge: Cambridge University Press.

Peirce, Charles Sanders (1910), 'Notes on the Doctrine of Chances'. In Justus Buchler, (ed.), *Philosophical Writings of Peirce*, New York: Dover, pp. 164–73.

Portner, Paul (2009), *Modality*. Oxford: Oxford University Press.

Schaffer, Jonathan (2007), 'Deterministic Chance?' *British Journal for the Philosophy of Science*, vol. 58: pp. 113–40.

Sober, Elliott (2010), 'Evolutionary Theory and the Reality of Macro Probabilities', in Ellery Eells and James H. Fetzer (eds.), *The Place of Probability in Science*, Dordrecht: Springer, pp. 133–61.

Stalnaker, Robert C., (1974) 'Pragmatic Presuppositions'. In *Context and Content*, Oxford: Oxford University Press, 1999, pp. 47–62.

——— (1998) 'On the Representation of Context'. In *Context and Content*, Oxford: Oxford University Press, 1999, pp. 96–113.

Thomason, Richmond H (unpublished, 2005), 'Ability, Action and Context', URL http://www.eecs.umich.edu/~rthomaso/documents/action/ability.pdf.

Williamson, Timothy (2000), *Knowledge and its Limits*. Oxford: Oxford University Press.

——— (2008), *The Philosophy of Philosophy*. Oxford: *Blackwell*.