Is the Past a Matter of Chance?*

Antony Eagle

August 21, 2012

One of the most elementary things we know about the connection between chance and time is that 'chances should only concern future events' Schaffer (2007: 124). Bigelow, Pargetter and Collins clarify what this feature of chance amounts to:

> Chance is temporally asymmetric. The future is chancy, but the past and present are not. If A is a proposition [wholly] about some time t, then the chance of A at some *later* time t' will take one of the extreme values one or zero, depending on whether or not A is true. (Bigelow *et al.*, 1993: 454)

Moreover, this has some claim to be a platitude about chance. It features, for example, in Lewis' discussion of the Principal Principle, elicited in part from his elementary questionnaire about beliefs concerning chance. He gives the example of the way that, as one negotiates a labyrinth, taking this turn and that, the chances of reaching the centre by noon vary. It is one of the 'firm and definite opinions' we have about chance that once one reaches the centre before noon, 'then and forevermore your chance of reaching it by noon is 100%' (Lewis, 1980: 91).

For clarity, here is the thesis I take Schaffer, Bigelow *et al.* and Lewis to be discussing (I use 'outcome' to avoid irrelevant issues about whether propositions or possible events are the basic chance bearers):

(☆) At every time, the only outcomes which are a matter of chance (that is, have non-trivial chances, strictly between zero and one), are future outcomes at that time. At any time when a given outcome is past, it is no longer a matter of chance (if it ever was), and its chance is trivial.

If it is indeed platitudinous, the claim that past outcomes are not a matter of chance may seem to stand in need of no further defence. However, recently

*Thanks to an audience at UMass Amherst, and to Alastair Wilson, Harjit Bhogal, and especially Chris Meacham for comments on a written draft.

arguments have been given that aim to undermine this piece of common knowledge about chance (Hoefer, 2007; 2011; Meacham, 2005). Such arguments don't show that (☆) isn't a platitude (even analytic truths can be doubted by those who understand them perfectly well (Williamson, 2008: ch. 4)), but they do deserve a response, and philosophical argument is called for when dogmatic assertion is not dialectically cooperative.

This paper aims to do two things. First, I'll explore the prospects of giving a positive argument for (☆). I investigate one popular argument in §1, and go on to give two better accounts: drawing on some recent work of Joyce (2007) in §2 I develop a preliminary argument for the general triviality of past chances. I give a deeper account in §3, explaining the results of §2 and giving further grounds for (☆), drawing on some of my own work (Eagle, 2011). Then, in §§4–5, I'll turn to evaluate some arguments against (☆). I conclude that our past is not a matter of chance, nor is the past a matter of chance in possibilities broadly like ours (in those that share the same physical laws, and the same sorts of causal dependencies). But I do concede something to the recent opponents of (☆), because there are possible circumstances in which the past (at those circumstances) is chancy. So it is not analytic of the concept of chance, nor metaphysically necessary, that the past is no longer chancy. But it remains true of, and something we know about, our actual chances.

1 Admissible Information and the Triviality of Past Chances

After noting that we ordinarily think of chance as time-dependent, Lewis offers an argument for the temporal asymmetry of chance in our sense. The argument rests on an application of the Principal Principle, Lewis' articulation of the connection between chance and credence. The Principle states that one's rational initial conditional credence in an outcome — conditional on the chance of that outcome in conjunction with any 'admissible' information — will equal the chance. Information is admissible, in turn, just in case it is not information that would trump the chances in its evidential bearing on the outcome — Lewis glosses admissible information as 'the sort of information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes' (Lewis, 1980: 92). The Principal Principle is intended to capture the way that chance is an *expert function*, guiding rational credence, for those epistemic agents who lack inadmissible evidence.

If chance is an expert for those who have only admissible information as evidence, then whether it is an expert for us, to which we should rationally defer our credence, turns on how much inadmissible evidence we have. We can use the Principal Principle to determine this, by noting that if ordinary people in ordinary situations do in fact rationally defer to the chances, then the kind of evidence they have must be the admissible sort. The questionnaire with which Lewis begins his paper establishes that 'a great deal is admissible' (Lewis, 1980: 92), because the rational opinion to have, in cases in which agents have the kind of knowledge we tend to have, conforms exactly to the predictions of the Principal Principle. Noting that the kind of information we tend to have includes historical information, Lewis proposes as one sufficient condition on admissibility at t that the information in question be 'entirely about matters of particular fact at times no later than t' (Lewis, 1980: 92).

With the Principal Principle and this sufficient condition on admissibility in our repertoire, here is Lewis' argument:

> Suppose the proposition A is about matters of particular fact at some moment or interval t_A, and suppose we are concerned with chance at time t. If t is later than t_A, then A is admissible at t. The Principal Principle applies with A [as admissible information]. If X is the proposition that the chance at t of A equals x, and if A and X are compatible, then
>
> $$1 = C(A|X \wedge A) = x.$$
>
> Put contrapositively, this means that if the chance at t of A, according to X, is anything but one, then A and X are incompatible. A implies that the chance at t of A, unless undefined, equals one. What's past is no longer chancy. The past, unlike the future, has no chance of being any other way than the way it actually is. (Lewis, 1980: 93)

The conclusion follows validly. But one might query the reliance on the supposed asymmetry of admissibility in the derivation of the temporal asymmetry of chance. If we focus on Lewis' original gloss on admissible information — information about outcomes that does not trump the information given by the chances — , rather than on the claim that being historical information is a sufficient condition for admissibility, then Lewis' argument can appear puzzling. Suppose we consider a fair coin flip yesterday; let A be the proposition that the coin landed heads, and let the actual outcome yesterday have been tails. Then here is an expression of some puzzlement about Lewis' argument:

> after the flip, [Lewis' sufficient condition for admissibility] would make $\neg A$ itself admissible; and of course $Cr(A|\neg AE)$ had better be zero.... But clearly this violates the correct definition of admissibility. $\neg A$ carries maximal information as to A's truth, and not by way of any information about A's objective chance; so it is inadmissible. (Hoefer, 2007: 554)

Hoefer makes this observation in the course of arguing that chance isn't time-dependent at all. His alternative explanation of the results of Lewis' questionnaire is not that historical information is admissible while information about the future is not. It is that we have asymmetrical *evidence* — there is a considerable amount of information about the past in our evidence, and relatively little information about the future in our evidence. And, says Hoefer, our evidence about the past contains lots of information about outcomes that trumps the time-independent chances, and is therefore inadmissible. The Principal Principle doesn't apply, conditional on a body of information including inadmissible information, so Lewis' argument misfires. But the eventual impact on credence is the same, since as long as credences are probabilities, where $E \vDash A$, $Cr(A|E) = 1$.

In response to this, of course, Lewis will simply deny that evidence about the past gives us information about past outcomes that is 'not by way of any information about objective chance'. That evidence *just is* information that the objective chances are now one or zero, and thus trivial. If we remain convinced of our prior chance assignments, and if Hoefer is right that those chance assignments are not time-dependent, then that should render our credences resilient to new admissible historical information. Since our credences are not resilient in the face of admissible historical information, the chances themselves must now be different than they were. (This is not to say that there is some mysterious change in the chance functions themselves. Rather, the expression 'the chance' is a description that, uttered at different times, denotes different probability functions, just as 'the prime minister' denotes different people.) For Lewis, the temporally variable denotation of 'the [present] chances' is mirrored by the temporally variable denotation of 'the admissible information' — for that picks out different bodies of information at different times. By contrast, for Hoefer, the chances are time-independent, and presumably the admissible information is time-independent too.

The upshot of all this, I think, is that there is little prospect of giving a suasive argument for the triviality of past chances from facts about which information is or is not admissible at a given time. The admissibility of information is too closely connected to the issue of which outcomes have trivial chances to be a helpful constraint on what value the chances at a given time have. Those who think the past isn't chancy, like Lewis, will and should take historical information to be admissible; those who think it is, like Hoefer, will and should take it to be inadmissible. And we simply do not have enough independent grip on the notion of admissibility to adjudicate this dispute. The dispute over admissibility simply recapitulates the dispute over chance, but without the resource of other principles that partially characterise admissibility — principles we do have in the case of chance. We will make more progress by addressing the chances of past outcomes directly rather than via admissibility of information about them.

❦

Another approach to this issue deserves some comment here. Rather than making use of the Principal Principle as an indirect connection between chance and credence, Handfield (2012) uses the credential role of chance to characterise directly what chance is in epistemic terms:

> The chance that P (at a time t, in a world w) is the degree of belief in P that is recommended by the best identifiable advice function, given only information that is available at t. (Handfield, 2012: 21)

To avoid triviality, this characterisation must be read as equating chance with the credence recommended by the best advice function that we are able to identify while making use of the information available to us.[1] Handfield (2012: 23–6) doesn't specify exactly which information is available, but it seems to be at least somewhat dependent on what kinds of information are appropriate in the context. Handfield and Wilson (forthcoming) are more explicit about the role of context in determining available information:

> A subject matter S is available in context C iff S is the most inclusive subject matter such that, for every (true) proposition E that is wholly about S, 'E could be known' is true in C. (Handfield and Wilson, forthcoming: §9)

Since, as noted above, there is considerable temporal asymmetry in the content of our evidence — an asymmetry which Handfield later explains on the basis of thermodynamic asymmetries (Handfield, 2012: §48; see also Albert, 2000: 113–25) — , this proposal will mean that all the functions we can identify as worthy advisors on the basis of our evidence are those that are recognisably expert with respect to the past. Since the best identifiable function will be amongst those we can identify as worthy advisors, the chance function will be recognisably expert with respect to the past — it will know all we know, and thus assign chance 1 to all past outcomes that we know.

So far so good for our claim (☆). However, consider a case in which we are ignorant of whether p is true, where p concerns some aspect of the past. In that case, some of the recognisably best advice functions recommend high credence in p, and some of them recommend low credence in p — to pick one of them as an advisor over any of the others would be to pretend to some evidence in favour of or against p that we simply do not possess. But then it will not be true that the best identifiable advice function assigns trivial chance to p, despite it being past. Indeed, in this case it looks like the definite description misfires, so there is no such thing as *the* best identifiable advice function. So the chance

[1] The alternative reading — which equates chance with the credence recommended by the best possible advice function when *it* is fed the information we have — is trivial: arguably the truth function Tr, which assigns 1 to all and only truths in its domain, is the best function; and if $Tr(p) = 1$, then $Tr(p|e) = 1$ for any available information e.

may be undefined; and even if defined, it's not clear that it will turn out to be trivial.

Handfield reduces the threat of undefinedness of chance somewhat, by being a bit more liberal about what information is available than this discussion has suggested. In Handfield (2012: 24) he draws a parallel with epistemic modals, and concludes that the available information is not limited to the information I have in evidence, but also information that there is a relevant way I could come to have in evidence after suitable investigation (see also DeRose, 1991: 593–4). In the revised account offered in Handfield and Wilson (forthcoming), the available information is that which is *knowable*, not just that which is known (i.e., in evidence, if E=K). Regardless of which liberalisation is more plausible, either narrows the field of identifiable advice functions somewhat, since now more information is available and fewer functions are expert with respect to all of that information. But unless we liberalise the notion of available information beyond all tenable connection with knowledge or epistemic possibility, and include all historical information, then there will inevitably remain in the set of potentially best advice functions some which disagree arbitrarily much over some historical fact about which we have irremediable ignorance.

The foregoing observations appear to show that (☆) is false if Handfield's account of chance is correct, because sometimes (in cases of irremediable ignorance of the past), the chances of past outcomes would not be trivial (though in many ordinary cases they would be). Again, however, things are not so clear cut. I'm not sure that Handfield's account is correct.[2] But even if the basic idea

[2]Here are a couple of reasons why:

1. Information about future outcomes can directly affect our current credence in those outcomes without altering our judgements of the chances. We can remain convinced that a yet to be tossed coin is fair, and has 0.5 chance of landing heads, even if we get information from a crystal ball that it will in fact land heads (which, since we actually have it, will be available information). But Handfield's proposal gives the wrong result in this case, because the best credence to have isn't 0.5 in this situation, in light of the available information. It's not plausible to insist in defence of Handfield that the chance would have to be trivial in this case, no more plausible than insisting that a proposition is necessary once there is information that it is actually true. Ultimately, I think the parallel case of modality is closely related, and that the blurring of epistemic and alethic modality illustrated by the last remark is akin to what Handfield attempts to do for chance. (Further remarks on this issue can be found in the exchange between Hoefer and Meacham, *this volume*.)

2. In conditions of ignorance about which advice function is best, one should set one's credence in p equal to one's subjective expectation of what the best advice is (that is precisely what the Principal Principle recommends in such a case, assuming that the chance function is the best advice function). Suppose the value of this expectation is x. There is no guarantee that the proposition *the right credence to have in* p *is* x has any credibility for you at all. (Suppose you think the best credence is either 1 or 0, and the available information is equally supportive of each hypothesis about what the best credence is. Then one's subjective expectation will be 0.5, and that is what one's credence should arguably be, even though it is not a credence that for you could turn out to be best.) It seems that either the chance is undefined, because there is no best identifiable function; or the chance is 0.5, since that is the right degree of belief to have, even though that is not — as it were — among the hypotheses about what it is best to believe one antecedently considers. Given the prior intelligibility

is right, attention to the details suggests that Handfield's account doesn't settle the issues we are addressing.

For Handfield's account to work, we need more constraints on the identifiability of an advice function than merely whether or not it agrees with our evidence on topics they both address. That simply leaves too many potential advice functions, and would make the chances undefined too frequently. (And it leaves in the wrong sort of advice functions — for of course the truth function will agree with our evidence on everything they both address.)

One way to impose a further constraint would be to say that the only candidate advice functions are those that would, in conjunction with the Principal Principle, determine the advice we would in fact take. That would certainly limit the field in the right way, leaving only the chance function in play. But it would also land us back in the situation of our earlier discussion, trying to figure out which information is admissible.

A more attractive constraint would result from appealing to other norms of rationality. Perhaps it could be a norm of rationality that *one should defer to the predictions of one's best confirmed scientific theories*. If so, and if those predictions

of hypotheses about chance, and the Principal Principle, we can conclude that the chance is not 0.5, despite the fact that the best credence to have, given the available information, is the subjective expectation of the chance, which is 0.5. (If we follow Handfield and conclude that the chance is 0.5, we can see this is self-undermining: the only reason we came to the conclusion that 0.5 was the best credence was by taking the subjective expectation of hypotheses about chance that we now know to be false and not deserving of any credence.) If so, this is a direct counterexample to Handfield's view. The direct response is presumably that it is not possible that one could know that the only live hypotheses about chance were that it was trivially one or zero. But it's hard to see what independent reason we could have for placing this restriction on what is doxastically possible.

3. Suppose we follow the later development of the ideas in Handfield and Wilson (forthcoming). One of the aims of the proposal is to undermine arguments for fundamentalism about chance, the idea that only the probabilities given by fundamental physical laws genuinely play the chance role, and that accordingly probabilities in games of chance, statistical mechanics, etc. are not genuine chances. Take the example of poker they discuss, where their account allows poker probabilities to be chances just in case no subject in poker contexts is be truly ascribed possible knowledge of propositions that are wholly about the actual card distribution. But of course this entails that poker probabilities are not chances — while it may be cheating to seize the deck and examine it, it is clearly a way of coming to know propositions at a level of grain finer than that appropriate for poker probabilities. I do not dispute that there is a good sense in which the rules of poker determine a relevant body of background information in light of which the chances are fixed (in earlier work of my own I advocate an account of this: Eagle 2011; see also §3). Rather, what I find hard to credit is that this body of information will constrain what it is possible to *know* — the principles governing inquiry (the rules of knowledge, about how to gather and consider evidence, etc.) are much less restrictive than the rules of poker. (There is nothing epistemically inappropriate about cheating at poker, so even if 'context determines that there are certain epistemic methods that are understood to be legitimate, possible, in use, or otherwise salient', it's false that 'when playing poker, the salient method … is one which does not involve violation of the rules of poker' (Handfield and Wilson, forthcoming: §9).) Drawing explicit attention to knowability ascriptions makes conversationally salient methods of inquiry that are able to rule out possibilities left open by the rules of poker alone. Handfield and Wilson may embrace this conclusion, but it's hard to see how doing so illuminates the functional role of probabilities in games of chance.

were in part probabilistic, then perhaps the probabilities of physics would be our candidate advice functions, and the best identifiable such function would be the probabilities featured in that physical theory which is best confirmed by the available evidence. While the truth function is perfectly good at making predictions, perhaps some further considerations mean that it cannot be a genuine candidate scientific theory.[3] If such considerations can be made out, we won't defer to the truth since the norms of rationality would forbid it being identifiable as a function which is worthy of deference (even though it is in fact worthy of deference). We don't have evidence sufficient to have knowledge in virtue of deferring our belief to the truth function, whereas we sometimes do have sufficient evidence to know what we believe on the basis of deferring to the chances.

Something about the foregoing must be correct. It is true that the predictions, probabilistic or otherwise, of our best supported scientific theories are worthy of special epistemic consideration. Perhaps the correct way to think about this is that scientific theories have a special status, and to then derive, in conjunction with Handfield's proposal, the claim that the probabilities in those theories deserve the honorific 'chance'. We would do better, it seems to me, if we begin by taking seriously the idea that scientific theories aim to give us the physical chances, and then use a truth about chance like the Principal Principle to establish that physical probabilities have a special epistemic role.

But on either approach, our need to limit the field of potential advice functions has plausibly left us only with physical probability functions as viable candidates. At this point, we need to figure out directly whether physical probabilities are the sort of thing which feature the kind of temporal asymmetry that would make the chances of past outcomes now trivial. (This is why I said Handfield's proposal is inconclusive.) In the next section, I turn to some observations about the characteristic features of physical probabilities that lend support to the idea that past outcomes generally have only trivial chances.

2 Deference, Resilience, and Physical Chance

Let us return to the idea mentioned above that chance is an *expert function* to which we ought defer our credences (Hall, 2004; Handfield, 2012; Joyce, 2007). As is now often observed, the Principal Principle, and the amended *New Principle* that Lewis (1994) offers to avoid some technical difficulties meshing the

[3]One such further consideration would derive from a Best Systems analysis of laws (Lewis, 1994). If the truth is not very simple, then the candidate scientific theories (systems of proposed laws of nature), will be so far ahead of truth in simplicity, and not so far behind in strength, that they are all much better than the truth function. The truth may be so complex that no scientist could endorse it at the end of inquiry; in that sense the truth may not even count as a scientific theory.

PP with his program of Humean Supervenience (concerning which, more below), is intended to capture the way that rational agents should defer their credences to a more expert probability function, namely, the chance function. Hall emphasises that the chance function is an expert *analyst* of any given information, in the sense that we should set our credences equal to $\mathrm{Ch}(p|e)$, where e is our (total) evidence (which includes our evidence about the actual value of the chances). The idea is that the chance of p, conditional on the information e, is the best opinion to have about p in light of e, and if e is our evidence, that will be the best opinion to have in light of our present evidence, which will be (unconditionally) the best opinion for us presently to have.

This conception of chance as an expert analyst gives no special role to unconditional chances at a time. For all this role requires, chance could be merely *sophistical* — skilled at analysing any given body of information, but indifferent about where the truth lies. But this fits poorly with how chances, at least those physical probabilities deriving from the natural sciences, actually guide credence. As Joyce observes,

> evidence about chances is somehow able to screen off information about the past — when neither *Cr* nor *e* contains evidence about the future, chance is not merely an analyst-expert, but an expert *tout court*. This is a key epistemological fact about chance, and it seems to be one of the things that differentiates chances from inductive or epistemic probabilities. When I know the indeterministic coin is fair or that the polonium atom has a half-life of 138.876 days, no amount of information about the past should lead me to shift my probability for the events in question (unless it first leads me to revise my view about the chances). The only way to explain this is to suppose that chance has some information at its disposal. Being a universal analyst-expert is not enough. When it comes to evidence about the past, chance must be a database-expert as well. What information does it have? One natural answer is to suppose that chance now knows everything there is to know about what is not now chancy. (Joyce, 2007: 199)

Joyce's proposal won't explain the temporal asymmetry of chance, since nothing in this argument either requires that the outcomes which are 'not now chancy' are past outcomes, nor precludes at least some of the presently chancy outcomes being past outcomes. Yet two important aspects of the epistemic role of chance are emphasised in this passage. First, chance is opinionated about the actual evidence, and cannot be indifferently conditionalised on any information whatsoever and still yield useful advice. In this, it contrasts with the usual way that evidential or epistemic probability is conceived, as giving something like an *a priori* probability of a claim conditional on any information whatsoever. (For Williamson, for example, the epistemic probability function gives 'the intrinsic plausibility of hypotheses prior to investigation' (Williamson, 2000: 211), and it should make sense to conditionalise this on any possible evidence whatsoever to yield the plausibility of hypotheses in light of that possible evidence.) It does not similarly make sense to conditionalise the chance function on any information whatsoever, at least not if one wants sensible results. For example, chances are *contingent*, if for no other reason than that they depend on which chance-entailing physical theories turn out to be true, and these physical theories are plausibly contingent. We will generally get poor results by conditionalising the actual chances on false information about the actual chances, even when chance is not expert with respect to chance itself; chance is not an analyst expert with respect to all subject matters. The actual chance function is actual because it embodies some contingently true information about actuality, which renders it unsuitable to be the chance function for every possibility, and correspondingly unsuitable to be an analyst expert without restriction. The contingency of chance is due to its being in part a repository of contingently true information.

The second aspect worth noting is that the particular kind of information chance has in its database often concerns the past. Indirect but compelling evidence for this is the fact that, as Joyce notes, *credences informed by known chances are resilient to the acquisition of evidence about the past*. New evidence about the past tends not to change the credences of agents who know the chances much, though it may change the credences of agents who do not know the chances significantly. Consider Alice and Bob, both weather forecasters. Both presently have credence 0.5 in tomorrow's being sunny — Alice because she knows the chance to be 0.5, and Bob because he knows the chance to be somewhere between 0 and 1, and has no information that disposes him to think it credible that it lies anywhere in particular in that interval. They both gain the previously unknown evidence that the last ten days were sunny. Alice's credence in tomorrow's being sunny does not change; she now knows that a somewhat unlikely series of events had previously occurred, but that doesn't make her any more or less confident in future outcomes. (If the evidence had been more misleading, of course, that might have undermined Alice's knowledge of the chances; but in this case it did not.) For Bob, however, this new historical information strikes him as evidence for some as yet unknown hypotheses about the chance, and evidence against other hypotheses. In particular, those hypotheses which make the chance of sunny weather greater than the chance of rain are supported by this new historical evidence at the expense of their rival contrary hypotheses. Bob no longer has a uniform distribution of credence over the possible hypotheses about chance, since he now assigns high credence to the chance of sun being greater than 0.5. And since he is, like Alice, reasonable, he conforms his credence to his subjective expectation of the chance, and comes to have a credence much higher than 0.5 in tomorrow's being a sunny day. His credence shifts considerably in light of that historical evidence, while Alice's does not — her credence is resilient. (Bob is unlucky here, of course, since his new credence is objectively more inaccurate than his old credence; he just happens to have received misleading evidence about the chances at a time when his opinion about the chances was vulnerable to such evidence.)

Thus described, Alice and Bob are both clearly rational in their very different responses to the historical evidence. The difference between them lies in Alice's antecedent knowledge of the chances. Her credence is resilient because the historical information, while not in general credentially independent of the future outcome (as Bob's response clearly illustrates), are conditionally credentially independent, given Alice's knowledge of the chances. Gaining historical information like *e* doesn't give Alice any different opinion about the chances of events she is concerned with. It is important to note in this case that the outcome is *not* in general probabilistically independent of the new evidence — the probability of a sunny day varies with respect to what the pattern of weather preceding it is. (Weather is unlike coin tossing.) So we cannot appeal to considerations of independence to explain the resilience of Alice's credence. The natural proposal at this point is to explain this by proposing that, in some relevant sense, the chance function which informs Alice's credence has already 'factored in' the preceding sunny weather, even if Alice herself was previously unaware of it.

So what is it that the chance function must factor in? The case of Alice prompts this proposal: *chance now is perfectly informed about the events that are causally relevant as of now to the presently chancy outcomes*. Joyce agrees:

> the present physical probability of an event encodes all information causally relevant to its occurrence that can be found in the present state…[4] (Joyce, 2007: 201)

The present chances screen off such information, rendering credences based on the chances resilient to such information, because any influence that the prior causes or present constitution can have over some possible outcomes is encompassed by the chances. There is no information about the prior causal history of some particular outcome (say, a fair coin landing heads) that would be evidentially relevant to that outcome to an agent who knew the chance of that outcome already — all the relevant causal information is conveniently summarised

[4]Joyce continues this sentence as follows: 'to know this probability is to know all there is to know about the present state as it pertains to the causes of the event'. This is not correct, since there is no reason why someone who knows the chances knows the causal story which grounds those chances. Compare: 'since the present facts about the water content of this glass encode all information relevant to its H_2O content, to know the water facts is to know all there is to know about the H_2O content'. This latter claim is false since agents needn't know the necessarily true but *a posteriori* reduction sentences that enable the derivation of the one body of information from the other. (See also fn. 7.)

by the chances.[5] It is particularly convenient since knowing the chances need not involve knowing all of this further detailed causal information; for limited agents like us, that makes chance particularly useful as a guide for objectively informed credence. This also provides the start of an explanation of the epistemic role of chance: we defer to the chances because, in summarising and accommodating a vast and heterogeneous body of information about past causes, it gives us vital information for prediction. Picking up on the contingency of chance, we might even propose to identify the chance function at a time as that probability function which is expert with respect to the causally relevant factors which have occurred by that time; the evident contingency of the latter will carry over to contingency in which probability function meets the description, *the chance function*.

I wish now to fend off one potential objection to the proposal of this section, as doing so enables me to clarify its scope and content. The objection begins with this observation: there are, very plausibly, cases in which the prior facts which fix the present chance of some outcome are not all causes of the outcome. One standard example is this, due to Schaffer (2000): suppose two independent spells are cast, one which has a 0.5 chance of causing $A \wedge B$, the other of which has a 0.5 chance of causing $B \wedge C$. As things turn out, A and B, but not C, occur. So only the first spell was a cause of B — the second spell did not succeed in causing B. Nevertheless, the second spell did succeed in raising the chance of B from 0.5 to 0.75. The supposed objection is that such cases show that a chance function can be perfectly informed about the past causes, but nevertheless fail to be perfectly informed about that on which the chance depends. This objection misfires: the proposal neither involves nor entails that being perfectly informed about prior causes of a given outcome is sufficient to fix the value of the chance, so pointing out another way in which such information is insufficient to ground the chance is neither here nor there. But seeing that this objection misfires does clarify my claim. I claim that the behaviour of chance — in its effects on credence, and its own contingent dependence on actual history — is best explained by suggesting that a constraint on which probability functions are eligible to be the chance function at a time is that those functions must be expert with respect to the causes of the outcomes to which chances are to be assigned. I do not claim that this is the only constraint on eligibility, nor that it is even close to being sufficient to fix the value of the chances. It is — at least so I argue below — only sufficient to fix the chances of those outcomes, information

[5]The resilience of the chance of some outcome to other information about the past is readily explained too, since other past information will be causally independent of the outcome and will be probabilistically independent too — since cases of probabilistic dependence without causal dependence involve common causes which will screen off those causally independent outcomes.

about the occurrence of which is included in the information about the extant causes of presently chancy outcomes; generally speaking, the past outcomes.

It is crucial to distinguish my proposal, that *the chance function is perfectly informed about the prior causes of the presently chancy outcomes*, from a proposal I reject, namely that *the chance function is perfectly informed about the factors on which the present chances depend*. The former proposal, unlike the latter, does not commit us to the anti-Humean thesis that none of the facts that *ground* the present chances are future facts. (The potential division over whether chance is perfectly informed about the causes of presently chancy outcomes versus being perfectly informed about the causes of the present chances is related to the dispute over whether present chances can be *undermined* by future outcomes (Bigelow *et al.*, 1993; Hall, 2004; Lewis, 1994).) Consider for example the most simple-minded Humean view of single case chance, *actual frequentism*: the implausible thesis that the chance of an outcome is numerically equal to the relative frequency of that outcome in the relevant class of trials. If chance was perfectly informed about what caused the chances to be what they are, then chance would be perfectly informed about every outcome, since the total pattern of outcomes is the supervenience base for the facts about relative frequencies. All chances, past and future, would be trivial if this were so. But the more moderate thesis I defend, that chance is informed, at a time, about the causes of outcomes that have occurred by that time, doesn't trivialise all chances. This example is merely illustrative: I do not endorse actual frequentism. Indeed, Joyce's observation gives us further reason to reject that view. For facts about relative frequencies are time-independent, and thus insensitive to causal information that we know, following Joyce, that the chances are sensitive to. Perhaps there is some more way of having an appropriate reference class for a given outcome be fixed in part by the time at which the relative frequency is to be ascribed. But it is difficult to come up with a plausible and specific proposal for picking out a contextually-appropriate reference class of outcomes that generates relative frequencies which mesh with our existing knowledge about chance (Hájek, 1997). (Though if the resources offered in §3 can be given a suitably Humean foundation, they might provide exactly the right sort of constraints for a contextualist solution to the reference class problem for some sort of sophisticated frequentist theory.) So much the worse for actual frequentism — but not for my proposal.

Given the foregoing discussion, I can now give an argument that the past chances are trivial. The chance of a given outcome varies in part because it is sensitive to the occurrence of events which are causally relevant to whether that outcome will occur. Since these occurrences happen at different times, chance varies over time, since at each time the chance of an outcome captures all the causally relevant factors for that outcome that have occurred by that time. That is, the chance at a time renders the outcome in question conditionally independent of these other occurrences. Since an outcome is obviously relevant to itself, the chance of any outcome which actually has occurred at a time should render the outcome conditionally independent of itself. The only way that can be is if the chance of that outcome is 1. So present outcomes (i.e., those that occur at the time their chance is evaluated) shouldn't be chancy. And since no future outcome is causally relevant to whether a present outcome will occur, no future outcome can subsequently alter the present chances,[6] so having been presently settled those outcomes remain settled from now on. The past is thus no longer a matter of chance.

This argument can be fleshed out a little more. Suppose that the laws of nature give us a specific contingent initial chance function, *Ch*. This might give chances of initial conditions (Albert, 2000), or those conditions may be independently specified. In any case, once we have some initial conditions, we have additional contingent information that needs to be taken into account in any prediction of subsequent events: that $Ch(p)$ has a certain value is one thing, but to rely solely on that would be to neglect other causally relevant information at the time of prediction. Since the chance function is supposed to give chances for all physically possible outcomes at a given time, the best way to accommodate this new physically contingent information at a time is to conditionalise the chance function on the causally relevant information at the time. So at each time, there will be a temporally indexed chance function $Ch_t(\cdot) = Ch(\cdot|R_t)$, where R_t captures information concerning the causally relevant factors, that have occurred by t, for any outcome in the domain of *Ch*. Since R_t will thus capture most if not all of the historical information at t, for most historical outcomes h (those which have occurred as of t), it will be that $h \in R_t$, and hence $Ch_t(h) = Ch(h|R_t) = 1$. To complete the argument that the chances of past outcomes are trivial, it suffices to argue that 'the chance of p', uttered at t, picks out $Ch_t(p)$ rather than $Ch(p)$. One route to this conclusion is via the deference-worthiness of chance; it is only if 'the chance' uttered at a time picks out a function which has factored in the causally relevant factors as of t that the credences set in line with the chances will be predictively optimal for an agent in that situation. Moreover, given that in general the information R_t will not be known, it cannot be plausibly argued that agents really do refer to *Ch* by 'the chances', and implicitly condition that prior chance function on the causal information.[7]

[6] Of course future outcomes could have played a role in grounding the value of the present chances, but not causally; and this grounding relation is time-independent, and so that couldn't subsequently alter the chances of an outcome that has already occurred.

[7] Just as agents needn't know what the present time is to use 'now' to pick out the present time — which shows that a present use of 'now' is not synonymous with August 21, 2012— , they

The preceding discussion suggests that the time dependence of chance is fundamentally explained by the expertise of chance with respect to the causal background of chancy outcomes, along with the *contingent* fact that information about the causes of actual outcomes is in fact all information about the past history of those outcomes. I do not propose here to attempt to explain the temporal asymmetry of causation, though I have discussed that topic briefly elsewhere (Eagle, 2007). I do not even claim that the direction of metaphysical explanation here is from the causal asymmetry to the asymmetry of chance. For all I've said, there is some independent reason for fixing on some body of information as appropriate for prediction, which would then indicate that the chance function should be expert with respect to that body of information. We might then propose some probability-raising account of causation (Glynn, 2011), and since all the information in I will have trivial chances at a given context, it will turn out that, at a that context, there are no prospective causes of outcomes in I; and we might then propose to use that package of views to explain the asymmetry of causation. I don't myself favour this approach; but since all my proposal needs is that there is an actual temporal asymmetry of chance, and that it is derived from the actual temporal asymmetry of causation, I see no reason to take sides on that further issue of reduction.

Finally, it is worth noting that the idea of a causal history can be given a relativistically acceptable gloss, so that the chance (at a point of spacetime) will be relativistically invariant. Chance isn't temporally dependent in any way that requires privileged foliation of spacetime.

3 Chance, Ability, and Context

The discussion in the preceding section explained the resilience of chance with respect to historical information by proposing that chance is a database expert about such information, which enabled us to derive the triviality of past chances. The argument is basically this: the behaviour of chance-informed credence is best explained by proposing that chance is expert about some information about causes, and I made a specific proposal about what information that might be, and gave an argument that led to the triviality of past chances on the basis of that proposal.

But one might wonder why chance has the feature of being resilient to historical information. That is the feature on which the inference to the best explanation of the previous section rested, and it would only strengthen that ar-

needn't know what the causally relevant factors are to use 'the chance' to pick out the present chance function, showing that 'the present chance of *p*' uttered at *t* isn't synonymous with $\mathbf{Ch}(p|R_t)$, even if the latter is the present referent of the former description.

gument if we had some reason for thinking chance must exhibit some database expertise. In this section, I will offer an argument that both bolsters the case that chance has this feature, and provides as well another route to the thesis that the past is no longer chancy. It will also cast the discussion of the previous section in something of a new light.

The argument of this section begins with consideration of another truth about chance, one I've defended at length elsewhere: that there is some chance of X φ-ing iff X can φ (Eagle, 2011); more crudely, that p has some chance iff it can be that p. The sense of 'can' involved in this claim is the dynamic (ability-attributing), circumstantial 'can'; it is not epistemic or deontic. This thesis provides a good way of making sense of the often noted modal aspect of chance. Leibniz argued that probability 'is a kind of graded possibility' (van Fraassen, 1980: 198), and this thesis tells us which sort of possibility corresponds to chances. It doesn't tell us how to get numerical values for chances out of facts about dynamic possibility — it's not a reduction of chance — but it does give substantive content to the anodyne Leibnizian thesis. It also generalises and explains the appeal of other claims about the connection between chance and possibility, such as the basic chance principle of Bigelow *et al.* (1993), without endorsing the more unattractive consequences of those claims. And most compellingly of all, the claim that chance is linked to ability ascriptions neatly explains the linguistic data, especially the unavoidable badness of claims like 'it can happen, but it has no chance of happening'.

If I'm right about this connection between chance and ability (as minimally encoded by the true 'can' claims), then we might expect the behaviour of dynamic possibility to be mirrored in the behaviour of chances. And this is what we do see, when we pay attention to the semantics of 'can'. The standard semantics for this expression, treated as a sentence modifier, says, more or less, that 'can p' is true at a context of utterance iff there is at least one metaphysically possible situation, amongst all those which circumstantially resemble the context closely enough (what is known as the *modal base*), in which p (Kratzer, 1991; Portner, 2009). How close is closely enough, and in what respect of resemblance? Those are a matter for context, which selects (somehow) a set of relevant features in view of which the ability is ascribed. There needn't be a principled story here about which features are and are not suitable to determine the modal base — we may need to look at actual usage to see how contexts actually function. In doing so, it is striking that we are almost without exception required to ensure that all worlds in the modal base agree on the intrinsic features of the subject of an ability ascription, and similarly required to ensure they agree on features external to the subject which would mask or defeat the ordinary effects of those intrinsic features. Ordinarily, when we evaluating the claim 'it can be that this cup breaks', we are required only to hold fixed its intrinsic features: 'It is in view of certain properties inherent in the cup, that it is possible that it breaks' (Kratzer, 1981: 64). But if the cup is saliently wrapped in bubble wrap, it would not be true to say that it can, in those circumstances, break (though it might yet remain *fragile*). There may be occasions when other factors are held constant, or when these default factors are allowed to vary; but in general it seems that these are the standard sorts of things that context picks up on.

This orthodox semantic story makes certain predictions about the behaviour of 'can' claims. In particular, it will predict that ability ascriptions are resilient to new causal information: given the factors which ordinarily make a direct semantic contribution to the truth of an ability claim (intrinsic properties of the bearer of the ability and the salient circumstances), the causes of those present properties and circumstances are rendered epistemically irrelevant to the question of whether the object will do what it can do. And we see this prediction borne out: 'can' claims display an intra-contextual resilience to new historical information: if a particular outcome is able to come about at a time, that modal claim remains true even given new information about the past events causally relevant to that outcome. If it's known that Clara can play the piano, the fact that she had piano lessons on July 8th, 1997 adds nothing to the case for her future piano playing, nor does the fact that she didn't manage to play last week because her dentist appointment ran long. Similarly, if it's known that Sylvester can't yet walk, information about his displaying last week all the symptoms of early walking won't prompt anyone to think that he might right now be walking. Again, the neatest hypothesis that explains this is that present information about an object's abilities conveniently summarises the causally relevant historical information pertaining to what is and is not possible presently for that object.

We now have one direct explanation for why chance-informed credence is resilient to new historical information, as in the case of Alice in §2: because chance-informed credence is a specific kind of ability-informed belief. Facts about abilities conveniently summarise the prior causally relevant factors as to whether the ability will be exercised, and so new historical information will have already been accounted for by the ability ascription in making informed judgements about whether the exercise of the ability will occur. And this is the same behaviour we see in chance-informed credence, though of course chances give us quantitative information about how likely it is that the exercise of the ability will occur. This gives us a deeper understanding of why chance exhibits the behaviour we took as our starting point in §2.

Attentive readers will be objecting at this point: *Hang on! It's perfectly pos-*

sible to know that an object has the ability to φ, and yet come across new historical information that is relevant to whether the ability will be exercised. Smith, rich and gravely ill, certainly has the ability to improve his children's financial situation. But if I then find out about his will, which bequeaths his entire fortune to the local animal shelter, I can successfully predict that he won't exercise this ability. So that newly uncovered historical information trumps the known ability, so the ability claim cannot conveniently summarise the prior causally relevant factors: the will is clearly relevant and yet excludes the exercise of the ability. This objection is astute, but depends on a false presupposition about ability ascriptions: namely, that if an ability is correctly ascribed against one context, it remains correctly ascribable in all contexts. The context-sensitivity of 'can' claims, emphasised throughout the literature, will predict that in many circumstances, the discovery of new historical information will make that information contextually salient, and alter the modal base by restricting it to those possibilities in which the newly discovered information holds. In the present case, the new information about the will is accommodated by restricting the modal base to worlds in which the will exists, and throughout the closest such worlds, where wills are enacted as they actually are, the children will not inherit anything. So there are no situations in the relevant space of possibilities in which the children inherit — in light of that new information, they cannot inherit. This phenomenon of accommodation of new historical information means that ability-ascriptions are contextually fragile, and that is why earlier I emphasised the *intra-contextual* resilience of ability claims to new information. Fixing a given body of contextually salient information, certain possibilities lie in the modal base for the 'can' operator; finding out more about the actual possibility won't undermine the proposition expressed in that context by a 'can'-involving ability ascription, though it may (through contextual accommodation) make impossible to express that proposition using the very same words in the new context. Despite the difficulty in using a 'can' claim to express it in another context, the ability truly ascribed by a 'can' claim in a context — namely, that a certain possible outcome is consistent with some contextually salient body of information — remains an ability in every context (consistency is not context-sensitive).

In every context, however, what an ability ascription picks out is a claim that summarises the contextually salient and causally relevant information with respect to the exercise of that ability. It doesn't seem to be the case in actual usage that which an ability can be relative to any old arbitrary body of information, but rather is ordinarily sensitive to just the kinds of causally relevant information about the intrinsic features of objects and the salient environmental masks of the ordinary activity of such features. The kind of trumping illustrated by the case of Smith above makes salient some information that we aren't used to treating as ordinarily relevant to the exercise of an ability, so there is some tendency to think it is not relevant to the ability, which explains whatever pull we might feel to the putative objection. But, with the correct context-sensitive semantics for 'can' in hand, we are able to see that the reasoning in the objection goes awry.

❦

Set aside for a moment the resilience of chances to new historical information. For the connection between chance and ability allows us to offer a direct argument that the past is no longer chancy. If we say, of some proposition p about a past event, 'can p', that utterance expresses a proposition that is true just in case the contextually salient information is consistent with p. Customary usage demonstrates that historical information is ordinarily contextually salient, since we are ordinarily extremely reluctant to assert that past outcomes *can be different than they in fact were*.[8] Since our past cannot be changed, it must be that the past is the way it was. Exploiting the connection between chance and 'can', we see that the past has no chance of being other than it was, and thus has chance one of being the way that it was. The past is no longer chancy.

That the past is not able to be altered may be because there simply exists a tendency for historical information to be contextually salient. But having just drawn attention to the role of contextually salient information, we may be able to offer an explanation of this fact. If we ask whether some past outcome can occur, although it did not — for example, 'can this coin land heads yesterday?'[9] — it is very difficult to avoid making the fact that it did not land heads yesterday contextually salient. Contextual accommodation can, it's true, generate sets of conversational background information that manage to make sense of very exotic utterances relatively smoothly, but it is extremely difficult to envisage an ordinary situation in which the question concerning the ability of a past outcome to actually be different than it is manages to express a genuine and non-trivial query. The question makes the actual past salient; given the salience of facts about the actual past, there are of course no possibilities in the modal base of this context in which the coin lands heads yesterday, since all

[8]Note the circumstantial 'can' involved is *not* the pure alethic modal 'possibly' — for it is possible that the coin landed heads yesterday, and indeed the coin *could have* landed heads yesterday. But those modals introduce a different, less constrained, modal base than circumstantial 'can', which is the modal of interest here.

[9]Once again, this question must be distinguished from 'Could this coin have landed heads yesterday?' To deploy, but only by way of analogy, a distinction familiar from the literature on conditionals, we might say that the 'subjunctive' modal invites us to consider alternative possibilities which need not hold fixed the historical fact that the coin did not land heads yesterday. By contrast, the 'indicative' question in the main text does hold fixed this fact; and the markedness of the phrasing of the question, it's near-ungrammaticality, suggests that it does so regardless of what we know or can discover about yesterday's state.

possibilities are consistent. This familiar behaviour makes it very difficult to avoid trivialising the ability of the past to be circumstantially different. Where p is a past outcome at t, then generally all worlds in the modal base for a 'can' claim uttered at t are alike with respect to p. There are extraordinary contexts where this default behaviour does not arise, such as contexts in which time travel and the prospect of backwards causation are salient (Lewis, 1976), but those are fairly atypical, despite their philosophical interest.

Accordingly, the past cannot be different than it now has turned out to be; so what can have happened is just what did happen. Given the connection between chance and ability, supported by the systematic similarities between 'can' and 'has some chance', this temporal asymmetry of abilities will be reflected in a temporal asymmetry of chances — the outcomes with some chance of happening are just those which did happen. Accordingly, any past outcome p which has failed to occur by now has no chance of happening, i.e., has chance zero. Correspondingly, the past outcomes which have occurred by now, $\neg p$, have chance one of happening, and we have our thesis (☆): past chances are trivial. This is entirely consistent, of course, with those chances having previously been non-trivial: the contextual salience of historical information at that context in no way involves that information being salient at earlier contexts. Information about chances can provide more fine-grained constraints on future outcomes than bare information about what can and cannot happen, but nevertheless the temporal asymmetry of those constraints is of a piece with the asymmetry of ability.

I've been at pains to emphasise the context-sensitivity of 'can' claims, and the corresponding context-sensitivity of chances. On some views of ability — for example, the view of Vetter (2010), which classifies abilities as a species of potentiality, conceived of as genuine properties that objects may possess (and not merely ersatz or thin as the properties predicated by true 'can' claims may seem) — the connection between context-sensitive 'can' claims and genuine abilities will be much more tenuous than on the account I've provided. They could still endorse the connection between 'has some chance' and 'can' that I defend; but it is more attractive, on such a view, to connect chance with the more metaphysically thick sorts of abilities that such views have in their ontology. Such is the proposal, I take it, that lies behind some versions of the propensity theory (Giere, 1973; Mellor, 1971): to link chances with dispositions or abilities, treated as legitimate ontological primitives (and not to be reduced to any facts of the Humean mosaic). Such accounts won't have the resources to explain the triviality of past chances in the way I have, though there is no reason they couldn't take it as a basic fact about propensities at a time that nothing has a propensity to produce past outcomes. Indeed, those who gloss propensities as something like causal tendencies (such as Giere) may appeal to the temporal asymmetry of causation to explain the asymmetry of propensity chances. I will return to the propensity theory at the end of §5; I mention it here to emphasise that my own view captures much that is attractive about the propensity theory, but is not committed to the metaphysics of propensities and is, I believe, explanatorily superior.

❦

These observations both support and complicate my discussion of the triviality of past chances. Clearly we have here a good deal of support for the claim that what was not has now no chance of having been, deriving from the claim that what was not cannot now have been, which has a natural and straightforward explanation using the standard semantics for 'can' and some plausible observations about contextually-determined features of the modal base. On the other hand, we have some examples, in cases of time travel and backwards causation, the successful treatment of which requires that the context not hold fixed some features that it ordinarily picks up by default (Eagle, 2011: §6). Suppose I toss a fair coin on stepping into a time machine, destination so far in the past that present outcomes are probabilistically independent of the outcome. The coin can land heads in the past, and can land tails; indeed, it has equal chance of both outcomes. This case, even if not actually realisable, is perfectly intelligible. To make sense of the truth of these 'can' claims, uttered in these possible contexts, we need the modal base to hold fixed the causally prior, but historically future, facts about the constitution of the coin and the circumstances of its tossing. We must permit the variability of the historically past, but causally future, facts about the actual outcome. And this will give us the result that, at least in some contexts, (☆) will be false: there are past outcomes which have a non-trivial chance. In discussing the case under consideration, this seems the right thing to say. If there are past outcomes which depend on present events in the same way as other chancy outcomes depend on present events, then while it might be that historical information that trumps the chances is more readily available than in the ordinary cases, there seems little to recommend the view that the epistemic impact of such information isn't trumping the chances but merely reflecting them. (Even if you wish to rest a defence of (☆) on the Joyce-inspired argument of §2, these cases of time travel, involving as they do backwards causation, will mean that the chance function should be expert about the future causes of past outcomes, and inexpert with respect to those outcomes themselves, so we get violations of (☆).)

I certainly think it is a strength of the present account that it successfully encompasses both ordinary circumstantial modality, as well as that which arises in deviant but still straightforward time travel cases, in the same elegant theory. And I think it gives the right verdicts in those cases. So there are possible contexts in which (☆) is false, because it entails falsehood sentences about what can happen. But we need to be careful in drawing conclusions from this. We've now detected some previously covert context-sensitivity in the expression 'has some chance', which means that (☆) is, unexpectedly, a context-sensitive claim, expressing different propositions at different contexts. Two points arising are worth discussing: (i) The fact that our central claim can express a falsehood in some contexts doesn't show that it expresses a falsehood in any actual contexts, or in those in which it was originally uttered; and (ii) there are dangers lurking for arguments against the claim which fail to attend to the context sensitivity involved.

Starting with point (i), it's important to note that (☆) as stated is not prefixed with 'necessarily', and I do not intend to claim that it is necessary truth, only that it is true. Given this, pointing out that there are contexts in which it expresses a falsehood is only germane if those contexts are plausibly actual. It is a contingent matter, in my view, whether there is time travel or backwards causation; however, it seems reasonable enough on present evidence to conclude that there are no actual cases of either. (This may be too hasty; I return to this issue in §4.) Even if there are actual time travel scenarios, and utterances of (☆) when such scenarios are salient may express falsehoods, that needn't impugn (☆), if the contexts in which it is originally proposed are not such contexts. There will be a danger of equivocation, but once that is pointed out, it is open to the defender of (☆) to say that the contexts they were expressing (☆) with respect to were ordinary, non-time-travel contexts, and as such (☆) expresses a truth when taken as intended. This response, crude though it may be, is effective against some of the main arguments that could be raised against the thesis, which themselves rely on no funny cases involving exotica like time travel and the like. If the two sorts of arguments above work in ordinary circumstances, then there is a successful response to those sorts of direct arguments that attempt to show that there are non-trivial past chances even in quite normal circumstances.

This brings us neatly to point (ii). The context sensitivity we've identified can pose problems for arguments against (☆). Consider, by way of example, this argument, a pastiche of themes from the literature on this topic. Assume determinism. If determinism is true, then the historical truths at any time entail (in conjunction with the laws) the historical truths at any other time. Assume the laws are admissible. So we may run this argument:

1. Either some historical information is admissible, or it is not.

2. The historical truths at any time entail in conjunction with the laws the historical truths at any other time.

3. If there is any admissible historical information at all, then all historical information at any time is admissible at every time. (from 2)

4. Applying Lewis' argument from §1, all chances are trivial — an outcome has chance 1 at some particular time iff the outcome occurs at any time. So these trivial chances aren't time-dependent. (from 3)

5. If historical information is inadmissible, then the only admissible information is the laws. In which case, the only permissible instances of the PP are: $Cr(p|L \wedge \ulcorner Ch(p) = x \urcorner) = x$. The only time it is ever rational to defer to chance is initially, since at all other times one is more expert than chance; so it is just as if chance is not time-dependent.

6. So chance is not time-dependent. (1, 4, 5)

The argument is supposed to establish that if determinism is true, then chance isn't time-dependent, and hence that there cannot be the variability with temporal position that (☆) requires. However, if the foregoing discussion is right, 'admissible' is a context-sensitive expression, since whether some information is admissible depends on which factors contribute to the modal base relative to which 'has some chance' is to be evaluated. Noticing this, we can see that the inference from (2) to (3) may be invalid: for (2) is a context-independent truth, while (3) is context-sensitive. If there are ever contexts in which the truth of the deterministic laws is not included in the modal base, then there can be admissible information but the outcomes needn't be trivialised, since there may be possibilities under consideration in which the laws are not held fixed. And, as I've argued elsewhere, there are such contexts (Eagle, 2011: §6). So this argument, plausible though it seems at first glance, can be resisted: even if determinism is true there can be temporal asymmetries of chance.

The upshot of this discussion is this. There is a link between chance and the causes of chancy outcomes, on both arguments I've presented. In odd cases, where the direction of causation is contrary to its actual direction, this means that past outcomes needn't have only trivial chances. This strikes me as the correct result, since those cases are certainly intelligible as cases where the past is a matter of chance. But those cases are not central enough to pose a problem for our thesis: indeed, as they are only problem cases if the links between causation, ability, and chance are as I've described, they vindicate (☆) in exactly the kinds of ordinary cases that we thought we were evaluating it with respect to all along.

Let's sum up the discussion of the past two sections. The context-sensitivity of chance ascriptions, and the consequent context-sensitivity of (☆), means that there is no one claim that every utterance of (☆) expresses. If a context-invariant expression of my view is wanted, it would be something like this:

(★) In most[10] contexts in which we make claims about chances at the actual world, the semantic value at that context of 'chance' will be a probability function which assigns only trivial probabilities to events earlier than the time of the context.

Revisiting the formalism introduced at the end of §2: each occurrence of 'chance' expresses a probability function $Ch(\cdot|R_c)$, where R_c is the contextually salient background information, and Ch is the initial probability function given by the laws. (I return to the question of whether Ch is, in any context, rightly called a *chance* function in the following section.) The discussion of the past two sections has indicated some constraints on what information is included in R_c. §2 argued that information about causally relevant factors for outcomes in the domain of the chance function that have occurred by the time of the context must be included, on the basis that the inclusion of such information was part of the simplest and best explanation of the resilience of chance to new historical information. This section has argued that we may understand the role of R_c as identifying the modal base for the circumstantial dynamic modal operator 'has some chance'. It also proposed that, because of the intimate connection between this operator and the modal 'can', there is further reason to suppose that the modal base for 'has some chance' includes historical information, because the inability of present chance devices to alter their past outcomes is equivalent to there being no present chance of those outcomes being different than they in fact were. Given these facts about R_c, it follows fairly quickly, via the sorts of arguments I've offered in these past two sections, that any information in R_c at a given context will be truly said to have only trivial chance at that context, and that (as it happens) most if not all historical information will turn out to be included in R_c. That gives us our invariant thesis (★).

4 Arguments from Physics for the Chancy Past

Our thesis would be false if it turned out that the correct conception of physical chance turned out to involve the actuality of a chancy past, given actual physics. This prospect is raised by some recent discussions about the nature of objective probabilities grounded in classical statistical mechanics, where some have argued that such probability measures play the chance role, as witnessed

[10]The reason for this hedge is discussed in the second part of §4.

by their role in prediction and explanation, but also assign non-trivial probabilities to many past outcomes (Loewer, 2001; Meacham, 2005). So far, there is no problem: as just argued, the possibility of a chancy past is consistent with all I have said. So long as the statistical mechanical probabilities provide information of the sort we know that chance does about the causes of the event to be predicted or explained (or about the structure of the space of possibilities over which the chances are defined), there will be no difficulty in reconciling the existence of non-trivial statistical mechanical chances of past events with ordinary judgements about chanciness. The potential difficulties would arise if those statistical mechanical chances were not connected with backwards causation, with ability, or with the kinds of dynamic possibility connected with them.

So the difficulties do not arise, for example, on at least some interpretations of quantum mechanics according to which there are chances of past outcomes (Price, 1994). The proposal is intended to avoid certain difficulties in reconciling quantum mechanics and special relativity. Some experiments, which agree very well with the predictions of quantum mechanics, seems to involve faster-than-light causal connections between distant bits of the experimental apparatus, violating special relativity. But, argues Price, if present properties of the measured system can be causally influenced by their future states, we can accommodate these results in an entirely relativistically acceptable way. On that view, there are genuine physical chances of outcomes, given posterior states. Usually, of course, these chances are trumped by the posterior time, since the experimental apparatus is small and is highly salient to the experimenters, so information about the past outcomes is generally already in the set of relevant factors, trumping the chances of those outcomes. But one can envisage a setup in which, while the post-outcome state is known, the prior state, which is a genuine probabilistic effect of the posterior cause, is not information which is available to trump the chances. Then the right credence to have about that past outcome is that informed by the physical chances, which are expert with respect to the causes of that outcome. This fits very well with the views about what chance is expert with respect to above: the factors which are causally prior to the chancy outcome. If those causes are sometimes temporally posterior to the effect, as proposed in this case, we would anticipate past chances, which is intuitively what we get on Price's theory.

However, in the cases of non-trivial statistical mechanical probabilities of past outcomes, we have more of a puzzle. For there is very little in the literature on statistical mechanical probability suggests the existence of backwards causation or dyanamical past possibilities in that theory; yet the probabilities in that theory bear the hallmarks of chances (Loewer, 2001; Meacham, 2005).

That said, it would need considerable argument to contend that statistical mechanical probabilities couldn't be linked to causal information in the way that has been suggested. And any such argument would be a plausible contender for an argument that the statistical mechanical probabilities were not, after all, chances, but some other kind of objectively constrained probability, with a more nuanced relationship to the genuine chances than that offered by Loewer or Meacham.

For example, statistical mechanical probabilities are constructed from the standard measure over the state space. Since the standard Liouville measure μ on state space need not be bounded (there are sets of states of arbitrarily large measure), we must normalise the measure to get values which formally look like probabilities. So if p is a statistical mechanical proposition — a set of states — and q is another, then the statistical mechanical probability of *p relative to q* is given by $\mu(p \cap q)/\mu(q)$, supposing all those sets to have well-defined measure, and $\mu(q) > 0$. Meacham (2005: 286) proposes that these relative probabilities, given 'their explanatory power and normative force', should be taken to be chances.

Doing so, argues Meacham, raises difficulties for our thesis. For suppose that $\mathbf{q}$ is a proposition partly about t_0 and partly about t_2, and we are considering the chance of $\mathbf{p}$, which is entirely about t_1, at t_2 (where $t_0 < t_1 < t_2$). There does exist the relative probability $P_{\mathbf{q}}(\mathbf{p}) = \mu(\mathbf{p} \cap \mathbf{q})/\mu(\mathbf{q})$; and this will in general not have a trivial value, since $\mathbf{q}$ is not informative enough to trivialise the relative chance. But then there is a non-trivial relative probability of $\mathbf{p}$ at t_2, after the time of the outcome. And indeed there may be a different non-trivial relative probability at t_2 relative to $\mathbf{r}$, where $\mathbf{r}$ is partly about t_0 and partly about $t_{1.5}$.

The difficulty for these proposed problem cases is understanding how to connect relative probabilities with chance simpliciter as mentioned in (☆); without some proposal for how to do this, these examples don't pose any challenge to our thesis. The problem cases work by permitting arbitrary propositions, so long as they have a well-defined non-zero Liouville measure, to have probabilities relativised to them. And this can be extremely useful — for example, if your evidence is one such proposition, then the statistical mechanical probability relative to your evidence is going to be a particularly useful objectively informed probability measure. Maybe it is even worth deferring to. But not every expert probability function is a chance function, so it is premature to conclude that these relative probabilities are chances.

One good reason for thinking they are not chances is that they are not resilient to new information about past causes. Consider the example above, and suppose that s entails q but is strictly more informative about the causes of p. Then $P_s(p) \neq P_q(p)$ in general. So $P_q(p)$ is not a database expert with respect to the causes or other enabling factors of p, and thus fails to fill part of the chance role as identified above in §§2–3. So not all relative probabilities are chances — indeed, it seems most plausible to suggest that the relative probabilities which are chances at a time are just those relative to information at that time about the prior causes of the outcome whose chance we're interested in (noting, of course, that the content of that information will depend on context in the way sketched in §3). That proposal only gives non-trivial chances for past outcomes if there are future causes of past outcomes, something we have little reason to believe in classical statistical mechanics.

Meacham's examples rest on two assumptions: (i) the identification of relative statistical mechanical probabilities with relative probabilities; and (ii) the claim that if q is a proposition partly about a time, the non-triviality of probability relative to q at that time entails the non-triviality of chance at that time. Laid out in this way, it's hard to see why we should accept claim (ii). Not every proposition q carries the sort of information that a genuine chance function does, the sort of information about causes identified in §§2–3. For the proposition q that, in a given context, carries the contextually salient information about the abilities of chance setups to yield outcomes, P_q is a relative probability that is the extension of 'chance' uttered in that context. But for no other proposition q' in that context will $P_{q'}$ be a probability function expressed by a use of 'chance' in that context. Meacham's example, given what he tells us about the proposition q, is not, in any context, an eligible referent of an occurrence of 'chance', and thus poses no threat to the claim that the past is not a matter of chance (as that sentence, when uttered in any context, expresses a truth).

That doesn't mean that Meacham's examples of relative probabilities cannot be understood in terms of chance; they simply may be explicitly *conditional* chances at a time.[11] Meacham says little about the time at which his relative probabilities are to be evaluated — that is because the facts about relative Liouville measure on state space are eternal facts. But the expression 'the chance' picks out different relative probability functions at different times, and of course the chance of outcomes can therefore vary from time to time. If we consider a time after an outcome has occurred, its chances will be trivial. But consider a time before a given outcome, say t_0 in our example above. We can easily understand the relative probability $P_{\mathbf{q}}(\mathbf{p})$ in terms of the chance of $\mathbf{p}$ at t_0, conditional on the information in $\mathbf{q}$. Consider, by way of analogy, conditional

[11] On my view all chances are obtained from conditionalising the probability Ch given by the laws on the contextual causal information R_c, but these are not in my sense conditional chances, because they can be picked out by an unconditional use of 'chance'. By saying that arbitrary relative probabilities may be treatable as conditional chances, I mean that they can be picked out by explicitly conditional constructions such as 'the chance of p given q', and only by such constructions.

chances in coin tossing cases. At a time before a coin is tossed, there is a 2/3 conditional chance of a sequence of fair coin tosses containing two Heads and one Tail, given that the sequence begins with a Head. This is unproblematic even though the event conditioned on occurs partly after the event whose conditional chance we're concerned with. Now, this proposal might not always work: even in Meacham's example, the chance at t_0 (which is not initial nomological probability) carries more information about t_0 and the prior history than $\mathbf{q}$ does. So it's not the case that all statistical mechanical relative probabilities will be conditional chances. That some of them are gives us enough, I think, to explain why Meacham's assumption that relative chances are chances is as appealing as it is, despite being incorrect.

The probability relative to some region of state space, and the chance conditional on that region, may come apart, if chance is expert with respect to facts lying outside that region. Unlike our coin toss case, we may be, and often are, ignorant of the chances in classical statistical mechanics, and have then to fall back on relative probabilities in light of our evidence as the best we can do. And that best may well in many cases be quite good enough. But being an objectively informed and good enough expert function isn't sufficient to be a chance function.

❦

One of the main reasons Meacham has for considering relative probabilities to be chances is that, he thinks, it is only by doing so that we can understand some otherwise troubling features of statistical mechanics. Under the standard phase space measure μ, almost all of the phase space is occupied by a region where entropy is high, and very little by the region where entropy is low. Entropy is, more or less, a property representing how disorderly a region of phase space is. This feature of the standard measure (together with various substantive but widely accepted further claims that I will not detour to consider—see Callender (2011: §2)), entails that relative to almost any proposition, the probability of high entropy propositions is much greater than low entropy ones. If we place no further constraints on statistical mechanical probability than those imposed by the measure on standard phase space, then we easily explain the second law of thermodynamics, that entropy always increases: taking a system moving along a trajectory through phase space that passes through a region of positive measure R, representing a particular macrostate, it is overwhelmingly likely that the system will subsequently end up in a macrostate of higher entropy (because high entropy macrostates are overwhelmingly probable relative to R). Unfortunately, the same reasoning entails that entropy increases towards the *past* as well: relative to any state of any given entropy, it is most probable that a trajectory passing through that state came from a higher entropy region, and will go to a higher entropy region.

Thus it seems, on this minimal version of statistical mechanics, that the relative probability of the initial conditions of our universe being like we think they are (some very low entropy state) is overwhelmingly unlikely. One could accept this quasi-skeptical conclusion, but the standard move in the literature is to argue that this conclusion is at odds with our evidence, and that this minimal statistical mechanics is to be revised. The revision standardly made is to impose, in addition to the standard phase space measure and dynamics, the *past hypothesis* (Albert, 2000):

> Earlier states do not have higher entropy than present states because we make the cosmological posit that the universe began in an extremely tiny section of its available phase space. (Callender, 2011: §2.1)

Relative to the past hypothesis, low entropy initial conditions have maximal probability, and (modulo worries about how to use a universal constraint on entropy to predict behaviour in a subsystem here on earth) this is supposed to explain why high-to-low entropy transitions are almost never seen, and low-to-high transitions are nearly invariably observed.

The puzzle for us is how to understand the dispute over the past hypothesis. Minimal statistical mechanics — without a past hypothesis — makes a probabilistic prediction, namely that there is a very low but non-zero probability of low entropy initial conditions. Statistical mechanics with a past hypothesis disputes this, saying that instead the low entropy past has a very high probability. We can sharpen this puzzle by thinking about how initial chance functions are related to relative probability in phase space. On the minimal view, the initial chance function $Ch_{\neg ph}$ is probability relative to the laws of statistical mechanics (which specify the permissible trajectories through phase space, and allow results about the synchronic measure over the space of possible states to translate into results about how likely various states are to be traversed over time). With the past hypothesis in place, particularly conceived of as a law itself, the initial chance function Ch_{ph} is phase space probability relative to the laws of statistical mechanics and the past hypothesis. But how could *either* of these have been *chances*, if the argument of this paper is correct: for haven't I just argued that only the future is chancy, and there is never a time when the initial conditions lie in the future?

Moreover, it is crucial that the correct initial nomic probability function be a chance function, since without that claim, some standard physical reasoning about this situation is deeply puzzling. The reasoning I'm thinking about looks like this: Assume that statistical mechanics with and without the past hypothesis are the two rival options. Assume also that our empirical experience is veridical, so we were in fact in a low entropy initial condition l. We may conclude either that something extremely improbable happened, since $Ch_{\neg ph}(l)$ is very small; or that something probable happened, since $Ch_{ph}(l)$ is very high. If these probability claims involve chances, we may use the standard methodological principle that *the occurrence of an outcome which had low chance according to T is disconfirmatory of T* to conclude that minimal statistical mechanics is false and that statistical mechanics with the past hypothesis is correct.[12] However, if these theories don't give initial chances, this methodological principle doesn't apply — and it is far from clear what to put in its stead that would yield the same conclusion. (Meacham (2005: 287–8) argues, persuasively to my mind, against one prominent alternative, that the probabilities involved are credences constrained by some application of a principle of indifference.) So unless the initial probabilities given by the laws are chances, we cannot explain some orthodox reasoning about physical probability.[13]

There is a tempting misunderstanding of the arguments I have offered that makes my position on this issue appear untenable. The discussion of §2 suggested that the candidate chance functions are those that can be obtained by conditionalising some initial probability distribution on some time-relative body of information R_t, and this seems to entail that all actual chances are time indexed, and assign trivial chances to outcomes that are the causal antecedents of the presently chancy outcomes, which will, at every time, include the initial conditions. This would entail that initial probability distributions are never chances, particularly not the distribution $Ch_{\neg ph}$ which assigns non-trivial probabilities to the actual initial conditions. And this in turn would give rise to the problem of accounting correctly for probabilistic explanation in statistical mechanics.

Thankfully, this is not a consequence of the proposal laid out above. Chance functions are not time-indexed, but indexed to certain bodies of information.

[12]This argument will be resisted by those who think the past hypothesis itself stands in no need of explanation, such as Callender (2004). For them, the right chance theory is $Ch_{\neg ph}$, and they will resist the conclusion that the truth of the past hypothesis disconfirms minimal statistical mechanics. The worry in that case is why accept that we are in a statistical mechanically abnormal universe, rather than a normal one with misleading evidence. But this issue is not settled by these considerations, so I suggest that the argument in the main text is only compelling for those who see the low entropy past as standing in need of explanation. If it does not, then the worry that initial probabilities won't be chances is less pressing in any case.

[13]Note that we certainly don't need either Ch_{ph} or $Ch_{\neg ph}$ to *now* be chance functions, in order to make reasonable inferences about which one *was* the initial chance function. Think about how the reasoning goes in a case I take to be analogous. Suppose I'm holding fixed some facts about coin tossing, and trying to infer whether a given coin is fair from the outcome sequence it has produced. I can certainly say: 'since the chance of the coin landing *HHHHHHHHH*... was so low if it were fair, and it did land that way, that's good reason for thinking it wasn't (and isn't) fair'. I use the probabilistic predictions of various hypotheses about what the chances *were* to give me reasons for favouring some hypotheses over others. There is no tension between noting that the present chance of an outcome is trivial, and using its differing earlier chances according to rival theories in combination with its occurrence to confirm one such theory at the expense of another.

That information, because it contains information about causes, generally happens to be temporally asymmetric, and this is what generates the temporal asymmetry which supports (☆). But not every possible contextually salient body of information has this much content. In particular, the empty body of information can, if context is right, be contextually salient, which ensures that 'chance' in such a context denotes the initial nomic probability function (or whichever such function is appropriate given the body of probability-involving laws salient in that context, since we may be discussing a false theory under a supposition rather than simply making claims about the chances).

Why can these initial probabilities given by the laws (or putative laws) be chances, even though they are the chances at no time? Think about the contexts in which we wish to make use of probabilistic claims. Very often we are making predictions about what will happen, at this is the kind of context I have been focussing on in the discussion in §§2–3. But sometimes we are offering probabilistic *explanations*, and there the relevant chances will often not be the present chances, but chances relative to some other body of causal information. Again, very often this will be the causal information relevant to some previous moment in time, and we will be explaining some past event in terms of the chance it had of occurring at some time prior to its occurrence. But the initial conditions had no causes, and no explanation can be offered of them except in terms of the laws. But if we are in a context where the initial conditions are the appropriate objects of explanation, and we attempt to understand why the initial conditions were as they were, the best we can do is defer to the relevant expert function; and the relevant expert is the chance function, holding fixed just the laws and no other actual information. In such a context, the expert function will be the initial probability function given by the laws, Ch. This is a highly eligible referent, because the null body of causal information is a simple and contextually available body of information. Of course Ch is readily trumped in other circumstances by one of the many Ch_{R_c} functions as the most appropriate referent of 'chance', but in these special sorts of contexts in which no contingent information about actual history is appropriate (presupposing as it does information that entails the initial conditions without explaining them), the initial probabilities are the most eligible referent of 'chance'. One way to get into such contexts involves contextual accommodation: if explanatory 'why' questions are asked that presuppose the unavailability of information salient in the context, at least sometimes that can remove information from the contextually salient background, and by doing so one can make salient chance functions that assign non-trivial probabilities to past outcomes, even the initial conditions, as appropriate for probabilistic explanation in such contexts. It is difficult to get into such contexts, and never appropriate to address predictive questions within them; the chances appropriate for prediction of future outcomes are all of the sort that trivialise the chances of past outcomes.

Earlier we saw that not all relative probabilities can be easily understood as conditional chances at times. But the initial chance function *Ch* given by the laws can be taken to lie behind all statistical mechanical relative probabilities, which are all conditional chances if the preceding argument is correct. But once again, while all relative probabilities are conditional chances, they are not all chances, in the sense that they are not all eligible referents for 'chance' in some context of use or other — because not all relative chances are obtained by conditionalising initial chance on the sort of information a chance function would be expert with respect to, in a given context.

5 Philosophical Objections to the Triviality of Past Chances

The main purely philosophical objections to our thesis have been forcefully made by Hoefer Hoefer, 2007; 2011. Here's how he phrases the challenge:

> Lewis claims, as do most propensity theorists, that the past is 'no longer chancy'. If *A* is the proposition that the coin I flipped at noon yesterday lands heads, then the objective chance of *A* is now either zero or one — depending on how the coin landed. (It landed tails.) Unless one is committed to the 'moving now' conception of time, and the associated view that the past is 'fixed' whereas the future is 'open' (as propensity theorists seem to be…), there is little reason to make chance a time-dependent fact in this way. I prefer the following way of speaking: my coin flip at noon yesterday was an instance of a chance setup with two possible outcomes, each having a definite objective chance. It was a chance event. The chance of heads was 1/2. So 1/2 is the objective chance of *A*. It still is; the coin flip is and always was a chance event. Being to the past of me-now does not alter that fact, though as it happens I now know *A* is false. (Hoefer, 2007: 554)

Hoefer's discussion raises this question: if we, as philosophically conservative B-theorists, deny the non-perspectival significance of the distinction between past and future, then mustn't we equally repudiate the objective significance of chance if it is dependent on this distinction? As argued in §§2–3, the reason the past at t isn't chancy at t — i.e., isn't assigned non-trivial chances by the chance function at t — is because events in the past of t are not able to be affected by what happens at or after t, or so we actually think. The pastness of past events in itself has no significance; it is the fact that such events are not susceptible to present causal influence that is significant. Chance is time dependent in a very mundane way — since the set of things picked out by the description 'outcomes which have occurred' is different at different times, and it is the outcomes which have occurred which are potentially causally relevant to the presently chancy outcomes, the present chances will change from time to time too. To propose a single probability function as the referent of 'the chance function' at every time is to undo the connections between chance, cause, ability, and prediction. The B theory does not entail anything concerning whether most descriptions are temporally rigid or not — it's perfectly B theoretically acceptable to think that 'the laws of nature' is a temporally rigid description, while 'the prime minister of Australia' is not. The argument that 'the chances' is temporally non-rigid that I have offered above, deriving that non-rigidity from the fact that which outcomes are causally relevant to a given chancy outcome varies from time to time, is completely B theoretically acceptable. We need only appeal to the B theoretic view of tense, since the causally relevant outcomes are, as a matter of actual fact (plausibly), a subset of those which have occurred. 'Past' and 'future' play a role, but we only need perspectival, B-theoretic accounts of pastness and futurity to make sense of time-dependent chance; we don't need A theoretic non-reductionism about The Present.

❦

Hoefer's main target is not our thesis, but the propensity interpretation of probability. He certainly thinks that one cannot be a respectable propensity theorist without endorsing the A theory:

> we can perhaps cobble together something meant to represent the propensity account, that is block universe compatible. But stripping away all the A-series linked metaphors leaves us with something that has lost so much of its intuitive content, it no longer clearly has anything to do with what we mean by 'objective probabilities' — that is, something about the world that deserves to guide expectations for future events, makes certain frequencies in outcomes more likely than others, perhaps even explains the frequencies we do see, and so on. (Hoefer, 2011: 86)

In the absence of the A theory, says Hoefer, there is no content that can be assigned to the notion of a propensity. There may be scope here for an indirect argument against (☆), if he can make the claim that it is only our residual, perhaps implicit, commitment to chance having propensity-like aspects that would lead one to endorse the time-dependence of chance. I don't think this argument works, since the only features of chance drawn upon in §§2–3 above were so central to the notion that an theory of chance which failed to capture them would be manifestly defective. Nor am I a friend of propensity theories (Eagle, 2004). But I think Hoefer's argument fails even if we endorse a propensity conception of chance: nothing in that theory is necessarily inconsistent with the B theory.

Consider the underlying metaphysics of the propensity theory, which Hoefer thinks requires the A-theory. That metaphysical picture is that there exist probabilistic dispositions, that ground chance ascriptions. It is not the deeply puzzling numerical aspect of propensities that Hoefer objects to, but rather the role of dispositions here, and in particular the distinction between manifested and unmanifested dispositions:

> If we think of a coin-flipping setup as having a propensity (of strength 1 / 2) to make events unfold a certain way (coin-lands-heads), then once that propensity has done its work, it is all over. The past is fixed, inert, and free of propensities (now that they have all 'sprung' and done their work, so to speak). (Hoefer, 2007: 554–5)

But it is not correct that dispositions require A theoretic distinctions between past and future. Take a simple disposition like 'fragility'. Suppose a vase is fragile at t, which explains in part why it breaks at t'. Of course the B-theorist eternalist should say that the disposition at t exists at t', but it won't be true that the vase is fragile at t' — it's already broken. So the way in which the dispositional property makes disposition ascriptions true can't be this simple minded proposal: if α's disposition to M when S exists at t, then 'α Ms when S' is true at t. Rather, we should have this more sophisticated proposal: if α's disposition to M when S is *instantiated* at t, then 'α Ms when S' is true at t. Every B theorist will need time-dependent property instantiation: otherwise they could not explain the dependence of total truth on the distribution of properties over the block universe (consider in this connection Lewis' doctrine of Humean Supervenience). The propensity theorist too may perfectly legitimately appeal to this notion. Let the propensity to land heads when tossed tonight be instantiated today but fail to be instantiated tomorrow, after the toss lands tails. Why isn't this enough of a change in the propensities to ensure the time-dependence of chances (since the chance at a time depends on the propensities instantiated at that time)? Part of this may have to do with the existence of closely related propensities, such as the generic propensity to land heads when tossed. Perhaps Hoefer thinks this is what propensity theorists accept, and thus he finds it possible for them to deny the chanciness of the past only if what's true (simpliciter) about the propensities changes over time. But this won't help, because this propensity won't change over time in the right way, so can't ever have been what the propensity theorist had in mind.

In sum, even if we falsely supposed that time-dependent chance can be motivated only given a propensity theory of chance, that would not require us to endorse the A theory. The time-dependence of chances is entirely neutral with respect to the A theory/B theory debate. I see no philosophical objection here to the picture of time-dependent chance articulated above. We have every reason to think that chance varies over time, and in particular, actually varies in such a way as to ensure that past outcomes are no longer a matter of chance.

Faculty of Philosophy, University of Oxford
Exeter College, Oxford
antony.eagle@philosophy.ox.ac.uk

References

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

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.

Callender, Craig (2004), 'There is No Puzzle about the Low Entropy Past'. In Chris Hitchcock (ed.), *Contemporary Debates in the Philosophy of Science*, Blackwell, chap. 12.

——— (2011), 'Thermodynamic Asymmetry in Time'. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, fall 2011 ed.

DeRose, Keith (1991), 'Epistemic Possibilities'. *Philosophical Review*, vol. 100: pp. 581–605.

Eagle, Antony (2004), 'Twenty-One Arguments Against Propensity Analyses of Probability'. *Erkenntnis*, vol. 60: pp. 371–416.

——— (2007), 'Pragmatic Causation'. In Huw Price and Richard Corry (eds.), *Causation, Physics and the Constitution of Reality: Russell's Republic Revisited*, Oxford: Oxford University Press, pp. 159–90.

——— (2011), 'Deterministic Chance'. *Noûs*, vol. 45: pp. 269–99.

Giere, Ronald N. (1973), 'Objective Single-Case Probabilities and the Foundations of Statistics'. In Patrick Suppes *et al.* (eds.), *Logic, Methodology and Philosophy of Science*, vol. 4, Amsterdam: North-Holland, pp. 467–83.

Glynn, Luke (2011), 'A Probabilistic Analysis of Causation'. *British Journal for the Philosophy of Science*, vol. 62: pp. 343–92.

Hájek, Alan (1997), ''Mises Redux'-Redux: Fifteen Arguments Against Finite Frequentism'. *Erkenntnis*, vol. 45: pp. 209–227.

Hall, Ned (2004), 'Two Mistakes About Credence and Chance'. In Frank Jackson and Graham Priest (eds.), *Lewisian Themes*, Oxford: Oxford University Press, pp. 94–112.

Handfield, Toby (2012), *A Philosophical Guide to Chance*. Cambridge: Cambridge University Press.

Handfield, Toby and Wilson, Alastair (forthcoming), 'Chance and Context'. In *this volume*.

Hoefer, Carl (2007), 'The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance'. *Mind*, vol. 116: pp. 549–96.

——— (2011), 'Time and Chance Propensities'. In Craig Callender (ed.), *The Oxford Handbook of Philosophy of Time*, Oxford: OUP, pp. 68–90.

Joyce, James M. (2007), 'Epistemic Deference: The Case of Chance'. *Proceedings of the Aristotelian Society*, vol. 107: pp. 187–206.

Kratzer, Angelika (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.

Lewis, David (1976), 'The Paradoxes of Time Travel'. *American Philosophical Quarterly*, vol. 13: pp. 145–52.

——— (1980), 'A Subjectivist's Guide to Objective Chance'. In *Philosophical Papers*, vol. 2, Oxford: Oxford University Press, pp. 83–132.

——— (1994), 'Humean Supervenience Debugged'. In *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, pp. 224–46.

Loewer, Barry (2001), 'Determinism and Chance'. *Studies in History and Philosophy of Modern Physics*, vol. 32: pp. 609–20.

Meacham, Christopher J. G. (2005), 'Three Proposals Regarding a Theory of Chance'. *Philosophical Perspectives*, vol. 19: pp. 281–307.

Mellor, D. H. (1971), *The Matter of Chance*. Cambridge: Cambridge University Press.

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

Price, Huw (1994), 'A Neglected Route to Realism about Quantum Mechanics'. *Mind*, vol. 103: pp. 303–36.

Schaffer, Jonathan (2000), 'Overlappings: Probability-Raising Without Causation'. *Australasian Journal of Philosophy*, vol. 78: pp. 40–46.

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

van Fraassen, Bas C. (1980), *The Scientific Image*. Oxford: Oxford University Press.

Vetter, Barbara (2010), *Potentiality and Possibility*. DPhil thesis, University of Oxford.

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

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