Chance and Context

Wilson, Alastair; Handfield , Toby

DOI:

[10.1093/acprof:oso/9780199673421.003.0001](https://doi.org/10.1093/acprof:oso/9780199673421.003.0001)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Wilson, A & Handfield , T 2014, Chance and Context. in A Wilson (ed.), *Chance and Temporal Asymmetry*. Oxford University Press, pp. 19-44. <https://doi.org/10.1093/acprof:oso/9780199673421.003.0001>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

19-44, Chance and Temporal Asymmetry edited by Alastair Wilson, 2014, reproduced by permission of Oxford University Press
<https://global.oup.com/academic/product/chance-and-temporal-asymmetry-9780199673421>

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Chance and Context

Toby Handfield & Alastair Wilson

ABSTRACT

The most familiar philosophical conception of objective chance renders determinism incompatible with non-trivial chances. This conception – associated in particular with the work of David Lewis – is not a good fit with our use of the word ‘chance’ and its cognates in ordinary discourse. In this paper we show how a generalized framework for chance can reconcile determinism with non-trivial chances, and provide for a more charitable interpretation of ordinary chance-talk. According to our proposal, variation in an admissible ‘evidence base’ generates a spectrum of different chance functions. Successive coarse-grainings of the evidence base generates a partial ordering of chance functions, with finer trumping coarser if known. We suggest that chance-attributions in ordinary discourse express different chance functions in different contexts, and we sketch a potential contextual mechanism for making particular chance functions salient. The mechanism involves the idea that admissible evidence is available evidence: propositions that could be known. A consequence is that attributions of objective chances inherit the relatively familiar context-sensitivity associated with the modal ‘could’. We show how this context-dependency undermines certain arguments for the incompatibility of chance with determinism.

1. *Introduction*
2. *Incompatibilism about chance*
3. *Lewisian chance and admissibility*
4. *The source of incompatibilism*
5. *Generalized chances*
6. *Two distinctions amongst chances*
7. *The coarse-grained chance spectrum*
8. *Constraints on coarse-grained chances*
9. *Which are our chances?*
10. *Against Lewisian chance*
11. *Conclusions*

1. Introduction

Making and testing chance¹ hypotheses is an important part of our epistemic practice, in everyday life as well as in science. However, the most popular conception of chance entails that, if fundamental physics is deterministic, then no propositions have non-trivial² chances. *A fortiori*, no ascriptions of non-trivial chances in the non-fundamental sciences are ever true. This incompatibilist conception of chance was endorsed by David Lewis ((Lewis, 1980), (Lewis, 1986)) and has recently been defended by Jonathan Schaffer (Schaffer, 2007).

In this paper we explore an alternative to the incompatibilist picture. A generalized framework for chance, emerging from recent work by Arntzenius, Loewer, Meacham, Hoefer and others, allows for a range of chance functions to be characterized through variation of a parameter which we call the *evidence base*. Contextual variation in what is included in the evidence base, we suggest, is essential for accounting for the semantics of chance-ascriptions.

Incompatibilist objective chance emerges as a special case of generalized chance: it corresponds to a maximally fine-grained evidence base. Doubtless incompatibilist objective chance is a special case of some philosophical interest. Nonetheless, we suggest that there is good reason to think that it is not the subject-matter of most ascriptions of objective probabilities made in ordinary contexts.

One of the central advantages of our proposal is that it allows for a more literal and charitable interpretation of ordinary discourse involving the word ‘chance’. Although Lewis treated ‘chance’ as a semi-technical term, he was clearly aiming to capture something close to our ordinary talk of chance. Chance-talk seems to be objective in nature, yet it also seems to be wholly innocent of fundamental physics. On our account, chancy phenomena populate many levels of nature, not merely the fundamental level; and gambling games really are ‘games of chance’, regardless of how physics should turn out.

¹ We will use ‘objective probability’, ‘objective chance’ and ‘chance’ interchangeably.

² By this we mean chances not equal to zero or one.

Our proposal is in a certain sense deflationary. We do not introduce any novel metaphysical apparatus to account for the compatibility of chances at different levels, or for our knowledge of and guidance by the chance norms. Levels are not taken to be *sui generis* entities superadded to the fundamental physics; they can be characterized naturalistically, and their importance to us can be naturalistically explained via features of the epistemic predicament we face as situated agents. The overall account is thus congenial to naturalistic and pragmatic attitudes towards metaphysics.

The theory of chance which most closely resembles ours is the ‘Humean’ theory set out by Carl Hoefer (Hoefer, 2007) which incorporates certain attractive generalizations of the Lewisian picture of objective chance. Like ours, Hoefer’s theory allows for events to correctly be ascribed different chances in different contexts, and thereby reconciles our knowledge of non-trivial chances with the epistemic possibility of determinism. However, Hoefer’s account is skewed in various ways by his commitment to an austere Humean metaphysics encompassing only actual regularities as truth-makers for chance-ascriptions. This limitation exposes his account to a worrying variant of the single-case problem for actual frequentism; it prevents him from taking chances to be intrinsic to chance set-ups; and it entails that objective chances can only be correctly ascribed to a restricted sub-algebra of propositions. Separating out Hoefer’s generalized characterization of chance from his controversial metaphysics will help to cast light on the semantics and epistemology of chance-ascriptions more generally.

The plan is as follows. In §2 we motivate the project by arguing against incompatibilism, and in §3 we give a brief exposition of the Lewisian framework for chance. §4 traces the source of Lewis’ incompatibilism to an assumption about admissibility of evidence, and considers what Lewis called the ‘reformulated principle’, which makes this assumption about admissibility explicit. In §5 we introduce a generalized version of the reformulated principle, and in §6 we show how different assumptions about admissibility generate a diverse range of generalized chance functions when combined with this principle. §7 characterizes generalized chance in terms of the Lewisian theory of subject-matters, and §8 shows how fine-grained chances trump coarse-grained chances. §9 proposes a contextual mechanism, revolving around the semantics of ‘could’, for selecting a

particular chance function: the central idea is that admissible evidence is available evidence. In §10 we argue that generalized chance, rather than Lewisian chance, is the subject-matter of our ordinary chance-talk. §11 is a conclusion.

2. Incompatibilism About Chance

In a postscript to ‘A Subjectivist’s Guide to Objective Chance’, Lewis emphatically rejected the possibility of reconciling chance and determinism:

If the chance [of Heads] is zero or one... then it cannot also be 50%. To the question of how chance can be reconciled with determinism, or to the question of how disparate chances can be reconciled with one another, my answer is: *it can’t be done*. (Lewis, 1986) p.118.

Incompatibilism is supported by the following line of thought. If determinism is true, then the current state of the world determines the complete history of the world. Then the chance that the world-history will turn out to be any way, *other* than how it is determined to be, is zero. And by corollary, the chance that the world-history will be the way it *is* determined to be, is one. Seemingly, this leaves no room for any non-trivial chances.

The compatibilist, in contrast, is moved by the seemingly undeniable fact that the chance of a well-made and fairly-tossed coin landing heads is 0.5 – or very close thereto. This opinion seems to be untouched by the epistemic possibility that the fundamental physical laws might be deterministic.

How can the compatibilist reply, then, to the seemingly powerful argument for incompatibilism? We suggest that the compatibilist should agree with much of what the incompatibilist says. Deterministic laws, given the current state of the world, are sufficient to determine that propositions about the future will have values zero or one. But the compatibilist should resist the claim that this excludes all non-trivial chance-ascriptions. There are ways of making true, non-trivial chance claims even in deterministic worlds.

What the compatibilist wants to deny is something like the following principle (named by analogy with the ‘causal exclusion principle’, familiar from discussions of physicalism):

Chance Exclusion Principle: If ‘the chance of φ is x ’ is true at t , then ‘the chance of φ is y ’ is true at t , iff $x = y$.

This principle guarantees that the trivial chances generated by deterministic laws³ exclude any additional non-trivial chances. We will take CEP to be the central commitment of the position – incompatibilism about chance – against which we intend to argue.

Given CEP, if fundamental physics is deterministic, then there are no non-trivial objective chances whatsoever. Even if fundamental physics is indeterministic, then CEP entails that any theories which are underpinned by deterministic theories (as statistical mechanics is underpinned by classical mechanics) cannot involve non-trivial objective chances.

This is all wrong. Whether determinism holds at some more fundamental level is *obviously irrelevant* to the uses of probability we make in higher-level sciences. There is endemic chaos between the macroscopic level and any level that could turn out to be fundamental. If the usefulness of chance-theorizing at the non-fundamental level depended on underlying non-trivial fundamental chances, then the usefulness of chance-theorizing would provide us with compelling evidence that fundamental physics is indeterministic. But we have no such evidence. So it is incumbent on us to seek an account of higher-level probabilities which accords with their epistemic accessibility and usefulness to creatures with only coarse-grained measurement ability, without making this accessibility and usefulness dependent on open questions in fundamental physics.

As well as sundering our metaphysics of chance from our use of objective probabilities in the special sciences, CEP threatens an error-theory of chance-ascriptions in ordinary language. When engaged in gambling games, for example, we make frequent ascriptions of non-trivial chances to particular outcomes.

³ Why think that deterministic laws do generate trivial chances? We assume that deterministic laws never have any chance of being violated: if the laws plus enough information wholly about how things stand at t jointly entail some proposition X , then the chance of X at t is 1. The assumption is intuitively plausible; it also follows from the combination of Lewis' proposal for the 'grounding argument' (see §4 below) with the Principal Principle. Thanks to Chris Meacham here.

According to incompatibilism about chance, we are not in a position correctly to make these assertions unless we know (or at least: have a justified belief) that determinism is false. This argument cannot just be assimilated to sceptical arguments more generally; determinism (even if not a necessary truth as Hume and Kant seem to have thought) is still almost universally regarded as an open scientific possibility and not as a sceptical scenario.

CEP is highly problematic. It leads to an unacceptable disconnection between our use of objective probabilities in science and our metaphysics of objective chance, and it threatens an error theory of chance ascriptions in ordinary language⁴. Why did Lewis adopt it? To begin to answer this question, in §3 we will set out some details of the Lewisian framework for objective chance.

3. Lewisian chances and admissibility.

The account of chance we offer is broadly in the ‘functionalist’ tradition, associated with David Miller, Hugh Mellor, Isaac Levi and David Lewis. Functionalist theories of chance recognise a sharp distinction between objective probability and subjective probability, characterizing the former in terms of the constraints it places on the latter. The principle that encodes this connection between objective and subjective probability is widely known, following Lewis, as the ‘Principal Principle’ (for short: ‘PP’). Expressing his commitment to functionalism, Lewis famously declared that the PP captures ‘all we know about chance’ ((Lewis, 1980) p.266). In the light of more recent developments of the functionalist project (for example: (Bigelow, Collins, & Pargetter, 1993), (Arntzenius & Hall, 2003), (Schaffer, 2003), (Schaffer, 2007)), this assessment may need to be revised. Nonetheless, everyone in the functionalist tradition takes the PP to be indispensable in characterizing chance, even if they deny that it is the whole story.

Here is Lewis’ first – and most widely-known - formulation of the PP:

⁴ Though one of us is sympathetic to an error theory of chance (Handfield, MS), the point is that *these* considerations are manifestly too weak to justify an error theory.

L1PP: Let C be any reasonable initial credence function. Let t be any time. Let x be any real number in the unit interval. Let X be the proposition that the chance, at time t , of A 's holding equals x . Let E be any proposition compatible with X that is admissible at time t . Then $C(A|XE) = x$. ((Lewis, 1980), p.266)

Expressed informally: any rational initial credence function, conditionalized on the chance of P and on any other evidence E which is *admissible*, delivers a credence in P that equals its chance.

L1PP makes essential appeal to a notion of admissible evidence in characterizing the normative force of chance. Lewis says the following about admissibility:

Admissible propositions are the sort of information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes. ((Lewis, 1980), p. 272)

This characterization states a condition for admissibility, without telling us which propositions meet this condition. But the consequences of L1PP depend crucially on the types of propositions that are counted as admissible. As Lewis observes of the principle:

‘If nothing is admissible it is vacuous. If everything is admissible it is inconsistent.’ ((Lewis, 1980) p. 272)

Getting clear on what sorts of evidence should be counted as admissible is one of the main purposes of this paper.

The characteristic feature of chance, as characterized by the PP, is that *knowledge of chances screens off admissible information*: if we know the chance of P , and if we conform to the PP, then our credence in P is independent of any other admissible information that we might know. Admissible information is information which bears on P only by bearing on the chance of P ; so knowledge of the chance of P renders admissible information epistemically redundant.

Seen in this light, it becomes clear that the existence of objective chances is a highly non-trivial matter. For a proposition P to have an objective chance at all, there must be a non-empty set of propositions S whose evidential bearing on P is

screened off by one particular proposition Q – the proposition specifying the objective chance of P . It is not at all obvious *a priori* that such sets must exist. Yet chance is pervasive in our epistemic practice. Any good theory of chance – be it Humean, non-Humean or Everettian – must have something to say about which features of reality are responsible for this fact. Here we wish to remain neutral between different substantive theories of chance, so we set these issues aside.

Nothing we have so far seen in the Lewisian theory entails CEP, or generates a conflict between non-trivial chances and determinism. Lewis' functional characterization of chance, in and of itself, is neutral on the question of compatibilism. But, as we have seen, he was nevertheless firmly committed to CEP. It is worth getting clear on exactly where incompatibilism enters into the Lewisian picture. That is the task of §4.

4. The source of incompatibilism

As was emphasized in the previous section, the constraints that the PP places on rational credence depend on the exact characterization of admissibility that is combined with it. Lewis claimed to give no 'definition of admissibility' ((Lewis, 1980) p.272), being 'content to suggest sufficient (or almost sufficient) conditions on admissibility' ((Lewis, 1980) p.272). These conditions are: *i*) that information entirely about the past is always admissible, and *ii*) that information about the dependence of chance on past history is always admissible.

One effect of these conditions is to impose a time-asymmetry onto chance: chances of propositions entirely about past history are always zero or one, while chances of propositions partly about the future can potentially be non-trivial. This time-asymmetry, according to Lewis, 'falls into place as part of our conception of the past as "fixed" and the future as "open" – whatever that may mean' ((Lewis, 1980) p.273). The conditions are also the source of the time-dependence of chance, which Lewis likewise took to be key to our ordinary conception.

Whether or not time-asymmetry is an essential aspect of the concept we deploy in ordinary chance-talk is a nice question; we return to it briefly in §7. But the real source of incompatibilism lies in the failure of these conditions on admissibility to place any restriction on the *degree of detail* of admissible information. Lewis is explicit about this lack of restriction:

Admissible information just before the toss of a coin, for example, includes the outcomes of all previous tosses of that coin and others like it. It also includes *every detail – no matter how hard it might be to discover* – of the structure of the coin, the tosser, other parts of the set-up, and even anything nearby that might somehow intervene. ((Lewis, 1980) p. 272; emphasis added)

The easiest way to see how this aspect of the Lewisian criteria on admissibility leads to incompatibilism is to look at the use that Lewis makes of the criteria. Despite officially regarding them only as suggestions and not as a definition, Lewis nonetheless integrates them into his theory of chance by combining them with L1PP to generate what he calls the ‘reformulated principle’:

Let C be any reasonable initial credence function. Then for any time t , world w , and proposition A in the domain of P , $P_{tw}(A) = C(A|H_{tw}T_w)$. In words: the chance distribution at a time and a world comes from any reasonable initial credence function by conditionalizing on the complete history of the world up to the time, together with the complete theory of chance for the world. ((Lewis, 1980) p.277)

Stripping this of world index to simplify, and adding a subscript to remind us that C is a rational *initial* credence function⁵, we have:

L2PP: For all rational initial credence functions C_0 : $P_t(A) = C_0(A|H_tT)$

Lewis talks as though L2PP is a harmless reformulation of L1PP, which ‘enjoys less direct intuitive support than the original formulation’ but ‘will prove easier to use’ ((Lewis, 1980) p.277). But L2PP is neither harmless nor a reformulation of L1PP; in fact, L2PP smuggles in three significant extra commitments.

The first commitment is the time-asymmetry of chance, as signalled by the time index on P and H . Further discussion of this commitment will be deferred to §9.

⁵ (Meacham, 2010) emphasizes the importance of this point.

The second commitment is an existence conjecture for the ‘complete theory of chance for the world’, T. Further discussion of this commitment will be deferred to §5.

The third commitment, and the one which is the true source of incompatibilism, derives from the ‘complete history of the world’, H. L2PP says that the chance distribution coincides with the result of conditionalizing any rational initial credence function on H and T. But in deterministic worlds, a complete description of the world at any time suffices to fix a complete description of the world at every time. A rational initial credence function conditionalized on: (i) the complete description of the world at some time and (ii) the laws, will deliver only trivial credences. Thus L2PP trivializes objective probability in deterministic worlds, whether or not the existence conjecture for theory T is correct and whether or not we buy into the time-asymmetry of chance.

We have traced the source of the incompatibilism in Lewis’ theory of chance to his assumption that arbitrary detail is always admissible in the historical proposition H. In the next section, we will look at the consequences of relaxing this assumption.

5. Generalized chances.

It is the presence of H_t – the complete history of the world up to time t – in L2PP that is responsible both for incompatibilism and for the time-dependence of chance. But, as we have noted, H_t was put in by hand during the move from L1PP to L2PP; it is not essential to the structure of the Lewisian framework for chance. This point has been recognised for some time; authors including Arntzenius (Arntzenius, 1995), Meacham ((Meacham, 2005), (Meacham, 2010)), and Hoefer (Hoefer, 2007) have shown how the functional role of chance captured by L1PP can be separated from the Lewisian criteria for admissibility, and how a principle of the same form as L2PP can be provided which is neutral with respect to time-dependence and with respect to incompatibilism. We will refer to this principle as the ‘Generalized Principal Principle’:

GPP: For all rational initial credence functions C_0 : $P_G(A) = C_0(A|G)$

The proposition G is what (Meacham, 2010) calls the ‘grounding argument’. It can be filled out in various different ways, generating various different chance functions P_G . Lewis stipulated G to be a conjunction of the world’s chance-theory (T) and the complete history of the world up to time t (H_t). In this essay we will retain the thought that G can be factored into a conjunction of chance-theory T and some factual information, conveyed by a proposition E . We can nonetheless characterize a range of chance-like functions by varying the *type* of proposition which is admissible as E .

In reformulating the PP, Lewis assumed (apparently without argument) that, for every combination of proposition A and history segment H_t , the ‘complete theory of chance’ T outputs a value in the $[1,0]$ interval. This is the second commitment mentioned in §4: it can be thought of as an existence conjecture for a true theory consisting of a giant conjunction of conditionals⁶, such that every possible history segment is related via some such conditional to a chance distribution. An immediate consequence of the conjecture is that the chance function is always a probability function defined over all contingent propositions.

It is hard to see what warrant Lewis takes himself to have for this existence conjecture. Nonetheless, it is easy to see how the conjecture generalizes beyond L2PP to GPP. For every combination of proposition A and appropriate contingent factual proposition E , the generalized conjecture says that:

$$T(A,E) \rightarrow [0,1]$$

In what follows we will refer to E as the *evidence base*. The evidence base is a proposition that conveys maximal admissible information about the world. Allowing different types of proposition to play the role of evidence base enables us to generate different generalized chance functions P_G ⁷.

⁶ These conditionals may be thought of on the model of counterfactual conditionals: the material conditional is inadequate to capture them. See (Lewis, 1980) p.275.

⁷ It is worth commenting on the relationship between the generalized framework we have set out and the debate on whether unconditional or conditional probability should be taken as the basic concept. We have a choice between taking conditional probabilities as primitive, and defining

Lewis suggested that T is a giant conjunction of history-to-chance conditionals of the form: ‘If the maximally-detailed history up to t were ... then the chance function would be ...’ In the generalized framework, T becomes a conjunction of evidence-base-to-chance conditionals: ‘If the evidence base were ... then the chance function would be ...’

The existence of T , for some given E , is a far-from-trivial matter. Lewis assumed that T exists for an evidence base corresponding to maximally-detailed past history; in what follows we will make the analogous assumption for certain other kinds of evidence base.

All of the groundwork is now in place to characterize a range of generalized chance functions by varying the evidence base. In the next section, we discuss some different choices for the evidence base, and the chance functions that they generate.

unconditional probabilities in terms of conditional probabilities, or taking unconditional probabilities as primitive, and defining conditional probabilities in terms of unconditional probabilities. Alan Hájek ((Hájek, 2003), (Hájek, 2007)) has persuasively argued that if we take the latter route we end up facing thorny problems with badly-behaved infinities, and we also encounter a particularly virulent strain of the reference class problem. Hájek urges us to take the lead of (Popper, 1956) in treating conditional probability as basic and unconditional probability as the defined notion. Lewis’ framework for chance, and the generalized framework set out in this section, incorporates this feature from the start. Although the individual P_G functions generated by a particular choice of grounding argument are functions from propositions to numbers in the unit interval, to specify an individual P_G we must specify G . And specifying G amounts to specifying a ‘chance setup’. Lewisian chance, which takes as grounding argument the conjunction of H_t and T , characterizes chance as the probability of a proposition conditional on the instantiation of one very big chance setup – the entire universe, specified in arbitrary detail up to a time. But different choices of grounding argument, as long as T is retained as a conjunct, will likewise give rise to chances conditional on the instantiation of various different kinds of chance setup. This inbuilt conditionality insulates the generalized Lewisian picture from Hájek’s concerns.

6. Two distinctions among chances.

Lewis took the evidence base E to be the entire history of the world up to the present, specified in complete detail⁸. It is instructive to see what sorts of generalized chances are obtained by varying this condition along a variety of dimensions. In this section we will discuss two cross-cutting distinctions amongst chance functions: time-dependent vs. time-independent chances and fine-grained vs. coarse-grained chances. It is unlikely that all of the generalized chance functions discussed in this section can in fact be referred to by the word ‘chance’ in ordinary discourse. However, distinguishing them will help to elucidate the generalized chance framework, and to demonstrate its flexibility.

We begin with the trivializing special case where everything is admissible. If we take the evidence base E to be the whole history of the world from the beginning to the end of time specified in complete detail, we obtain chances which:

- ❖ Do not change over time.
- ❖ Are trivialized for all propositions.
- ❖ Are trivialized whether or not fundamental physics is deterministic.

Call this chance function the *omniscient function*. The omniscient function takes a proposition about any subject matter, and returns the value 1 if it is (timelessly) true and 0 if it is (timelessly) false. If we set our degrees of belief according to this function, we will have credence 1 in all the truths and credence 0 in all the falsehoods. A PP for the omniscient function is a simple expression of accuracy as the goal of rational belief.

There can be no doubt about the existence conjecture for T in the context of the omniscient function. As long as an algebra of propositions is closed under classical negation and conjunction, the omniscient function will be a probability

⁸ In the light of special relativity, it would be natural to replace this condition with the condition that E be a description of the contents of the past light cone, specified in complete detail. This complication will be irrelevant to our discussion.

function. (Note that there can be no such thing as the notorious ‘problem of undermining’ or the ‘big bad bug’⁹ for timeless chances.)

Of course, we don't (and can't) know which function, among the myriad of possible assignments of truth and falsity to propositions, really is the omniscient function. The advice that a PP for the omniscient function offers is so general as to be useless in practice; this is why Lewis referred to it as ‘vacuous’. However, the omniscient function occupies one end of a interesting spectrum of chances; and bearing it in mind can help us get clear on the role that other chance functions – which may be less informative, but easier to discover – play in our epistemic lives.

One such less-informative chance function is the Lewisian chance function. If we take the evidence base E to be the history of the world specified in complete detail up to the present, we obtain chances which:

- ❖ Change over time.
- ❖ Are trivialized for propositions about the past.
- ❖ Are trivialized if fundamental physics is deterministic.

The Lewisian evidence base includes only past history, rather than total history; however, it still includes all the true propositions about that history, no matter how detailed. Since what is past changes over time, so does the evidence base, and so does the Lewisian chance function it generates.

Note that it is no longer obvious, once we move from the omniscient function to the Lewisian chance function, that there will still exist a suitable set of true conditionals to serve as the chance-theory T. As mentioned above, Lewis seems to have assumed without argument that there exists a conjunction of evidence-base-to-chance conditionals T which will generate a probability function defined over (nearly) all propositions, and whose probabilities will screen off all other admissible information.

⁹ See (Lewis, 1986), (Lewis, 1994), (Hall, 1994), (Ismael, 2008) and (Briggs, 2009) for a sample of the literature on this problem, which results from combining Lewis’ functionalist theory of chance with his Humean metaphysic.

(Lewis, 1980) maps out the properties of Lewisian chances and argues that they capture many of our ordinary beliefs about objective probability. We think that Lewis overestimates the usefulness of Lewisian chance in accounting for our ordinary beliefs expressed using ‘chance’ and cognate terms. In §2 we discussed the unwanted incompatibilist consequences of Lewisian chances, and in §4 we traced these consequences to Lewis’ particular choice of evidence base. We accordingly turn next to a type of generalized chance function which is more germane to our compatibilist project.

Having criticized the Lewisian choice of evidence base, what alternative do we have to offer? Our proposal is to nominate for the evidence base propositions with a more *coarse-grained subject matter*, in a sense to be elaborated below. For simplicity we discuss first the timeless case. If we take the evidence base E to be a *less finely detailed* specification of the whole history of the world, we obtain chances which:

- ❖ Do not change over time.
- ❖ Are trivialized only for sufficiently coarse-grained propositions.
- ❖ Need not be trivialized, even if fundamental physics is deterministic.

We will make the notion of fineness of grain of description more precise in the next section. But an intuitive grasp of the notion suffices to understand the nature of coarse-grained chance. Coarse-grained chances are objective probabilities conditional on all the facts from a particular domain – for example, all of the facts about the thermodynamic macro-history of a world. Of course, there are many different ways of describing a history in less specific detail, so there are many different kinds of coarse-grained chance.

As with Lewisian chances, there is no *a priori* warrant for the existence conjecture for the chance-theory T for coarse-grained timeless chances. But in particular cases we may find *a posteriori* reason to believe the T existence-conjecture. The Liouville measure in statistical mechanics appears to be a timeless coarse-grained chance function; and the equilibrium distribution of particles in

Bohmian mechanics plays a similar theoretical role¹⁰. In what follows we will set aside the delicate question of how these conjectures are to be empirically confirmed and instead focus on their consequences.

Timeless coarse-grained chances find a natural home in theorizing about statistical mechanics and about Bohmian mechanics. But in more ordinary contexts, we may have a use for coarse-grained chance functions which are time-dependent in the same way as Lewisian chances. If we take the evidence base *S* to be the history of the world specified in less fine detail up to the present, we obtain chances which:

- ❖ Change over time.
- ❖ Are trivialized only for sufficiently coarse-grained propositions about the past.
- ❖ Are not trivialized if fundamental physics is deterministic.

Examples of coarse-grained time-varying chance could be the chance of rain on your next birthday, conditional on the current macrostate of the world; or the chance of winning a poker hand on which you and your opponent have both gone all-in, conditional on the cards currently visible.

In the next section we will characterize coarse-grained chances more precisely, and give a sketch of how different coarse-grained chance functions relate to one another.

7. The coarse-grained chance spectrum

As we reduce the fineness of grain of information going into the evidence base *E*, so as to include successively fewer details, so the nature of the resulting chance function undergoes successive changes. Saying more about these changes will require a more precise statement of what we mean by a coarse-graining.

The idea of one subject matter being more coarse-grained than another is roughly that one subject matter can be more or less *informative* than another.

¹⁰ See (Loewer, 2001) for more discussion.

Propositions about a coarse-grained subject matter do not make very fine distinctions between possibilities, compared to propositions regarding a more fine-grained subject matter. So by knowing something about a fine-grained subject matter, one has more information than if one knew something about a coarse-grained subject matter. The natural way to understand this property of ‘informativeness’ is in terms of entailment, or inclusion relations. Taking propositions to be sets of worlds, one proposition, P , is more informative than another, Q , if P is a proper subset of Q (i.e. P entails Q).

Lewis (Lewis, 1988b) (Lewis, 1988a) helpfully shows how these notions of relative informativeness can be rigorously extended to apply, not just to individual propositions, but to entire classes of proposition – or subject matters. Here are the essential features of the theory. A subject-matter is a partition of (equivalently: is an equivalence relation on) the set of all possible worlds. Intuitively, we can think of this partition as dividing worlds into classes which are equivalent with respect to the subject-matter constituted by the partition. So if our subject-matter is the 17th century, then all the worlds in any cell of the partition have duplicate 17th centuries.

A partition A is a subject matter of a proposition P if and only if P is a union of cells in A . Or equivalently: if and only if each cell of A either implies or contradicts P . Trivially, every proposition P has as a subject matter, *whether or not* P : the partition of worlds into P and not- P . But there may be other, less trivial subject matters associated with a proposition. ‘There are seven coloured objects on the mat’ has the subject matters: how things are on the mat; where those seven objects are located; how coloured objects are distributed; and so on.

What is it for one subject matter B to be a coarse-graining of another, A ? It is for B to be *included* in A . Reverting to our first example, the history of the 1760s in Scotland is included in the history of the 17th century. The topic of that decade in Scotland is a coarser subject matter, because it ignores various facts that are properly part of the history of the 17th century. Subject-matter A includes subject-matter B iff all worlds which are alike with respect to A are alike with respect to B : that is, iff no worlds that A places into the same cell are placed into different cells by B . Or equivalently, iff all cells of B are unions of cells in A . The most inclusive

subject matter, then, is simply the partition such that every cell contains only one world – this is the degenerate subject matter, *how everything is*.

The main attraction of the Lewisian theory is that inclusion of subject-matters is just the well-understood notion of supervenience. (The facts about Scotland in 1760 supervene on the facts about the 17th century in general.) Since supervenience generates a partial ordering on partitions of worlds, inclusion generates a partial ordering on subject-matters. And since inclusion generates a partial-ordering on subject-matters, our coarse-graining relation generates a partial ordering on evidence bases.

Let us unpack this a bit by giving examples to illustrate the formal properties of coarse-graining as a partial ordering of evidence bases.

Coarse-graining is transitive: if A is a coarse-graining of B, and B is a coarse-graining of C, then A is a coarse-graining of C. A simple example: A is a specification of which objects are coloured; B is a specification of which objects are red, and which objects are coloured; C is a specification of which objects are scarlet, which are crimson, which are red, and which are coloured.

Coarse-graining isn't euclidean: A and B can both be coarse-grainings of C, without either A being a coarse-graining of B or B being a coarse-graining of A. A simple example: A is a specification of which objects are coloured; B is a specification of which objects are red and which objects are coloured; C is a specification of which objects are green and which objects are coloured.

Coarse-graining is reflexive: everything supervenes on itself, so every subject-matter includes itself and every evidence base is a coarse-graining of itself. We could, if we wish, speak of proper coarse-grainings, by imposing the constraint of non-identity as well as the constraint of supervenience. We will have no need of such a notion.

This concludes our summary of the Lewisian theory of subject-matters. The most important message to take away from it is that the relation of inclusion amongst subject-matters generates a partial ordering of evidence base-types according to their qualitative informativeness. This partial ordering of evidence base-types generates a corresponding partial ordering of coarse-grained chance

functions. For every coarse-grained chance function, we obtain a coarse-grained PP expressing the ‘chance norm’ for that function. Our framework therefore provides a partial ordering of chance norms.

The discussion thus far has been rather abstract. In the remainder of this section, we will discuss three potential applications of coarse-grained chance, each of which is generated by a coarse-grained evidence base-type, and each of which corresponds to a distinctive chance norm.

The first application is Bohmian mechanics (BM): a deterministic physical theory, popular amongst philosophers of physics, which was developed in response to foundational problems in quantum mechanics. Although BM is deterministic at the underlying level, it nonetheless gives rise to effective indeterminism due to its inclusion of a set of ‘hidden variables’, on which the future evolution of the system sensitively depends, but which (assuming an initial ‘equilibrium distribution’) cannot be measured with unlimited accuracy. BM thereby generates a non-trivial conditional chance function if we set the evidence base to exclude information about the hidden variables; this chance function is (again assuming an initial equilibrium distribution) given by the same recipe as are the Born rule probabilities in ordinary quantum mechanics.

The chance norm for the Bohmian chance function can be stated as follows (where E_{BOHM} consists in evidence about the wavefunction of a quantum system, but not about the hidden variables):

$$\mathbf{PP}_{\text{BOHM}}: \quad P_{\text{BOHM}}(A) = C_0(A|E_{\text{BOHM}}T)$$

The second application is to statistical mechanics (SM), a theory with deterministic underpinnings which is capable of explaining a very significant number of features of thermodynamics. Like BM, SM involves ‘hidden variables’. In SM, the hidden variables are the ‘microscopic’ features of a system: typically, the positions and momenta of the system’s constituting molecules. The ‘macroscopic’ features of a system – temperature, pressure, etc – are easier to measure but do not uniquely determine the microscopic features.

The chance norm for the SM chance function can be stated as follows (where E_{SM} consists in evidence about the macrostate of a thermodynamic system.)

$$\mathbf{PP}_{\text{SM}}: \quad P_{\text{SM}}(A) = C_0(A|E_{\text{SM}}T)$$

SM incorporates a specific measure over the space of allowed histories – the Liouville measure – which is empirically successful and has a number of theoretically attractive features. This measure can be used to generate a SM-specific component of the global chance-theory T , providing partial *a posteriori* support for the existence conjecture.

The third application is to probability in games of chance; we will consider a simplified version of poker with no discarding of cards, no betting, and no folding. Again, we can isolate ‘hidden variables’ – the actual order of the cards in the deck prior to dealing, which determines the hands that will be dealt to each player. The winner of any round of this simplified game is determined in advance by the order of cards in the deck; so as with BM and SM, we have a deterministic process at the finer-grained level.

Despite the underlying determinism, the rules of the game generate effective indeterminism. The initial symmetry of our evidence between all possible orderings of cards in the deck allows us to calculate the likelihood of winning based only on the visible cards; indeed, these chances are so well-understood and well-behaved that (so we are informed) they are shown in real-time to viewers of televised poker tournaments. The chance norm for the poker chance function can be stated as follows (where E_{POK} consists in evidence about the cards visible to the player according to the rules):

$$\mathbf{PP}_{\text{POK}}: \quad P_{\text{POK}}(A) = C_0(A|E_{\text{POK}}T)$$

How do the different chance norms interact with one another? Suppose that SM is the correct physical theory, and that a poker game is in progress. If you know the microstate, you are in a position to know who will win the poker hand. Because E_{SM} entails E_{POK} , knowledge of P_{SM} trumps knowledge of P_{POK} . Knowledge of coarser-grained chance is made redundant by knowledge of finer-grained chance, because the latter take everything into account that the former does (and more.)

In the next section, we'll see in a little more detail how this works.

8. Formal features of coarse-grained chance

The chance role, as characterized via the PP, places significant constraints on the relations between different coarse-grained chance functions. We assume that the following constraint holds of any two genuine chance functions P_C and P_F , where evidence base E_F is a fine-graining of evidence base E_C :

Alignment: $P_C(A|E_F) = P_F(A)$

Alignment ensures that the chances at the coarser and finer levels harmonize with one another: PP_F and PP_C will never give conflicting advice. We can prove this by deriving PP_F from PP_C and Alignment, as follows¹¹:

1. $P_C(A|E_F) = P_F(A)$ (Alignment)
2. $P_C(A) = C_0(A|E_C T)$ (PP_C)
3. $P_C(A|E_F) = P_C(AE_F) / P_C(E_F)$ (Probability calculus)
4. $P_C(A|E_F) = C_0(AE_F|E_C T) / C_0(E_F|E_C T)$ (2, 3)
5. $P_C(A|E_F) = C_0(AE_F E_C T|E_F T)$ (4)
6. $P_C(A|E_F) = C_0(AE_F T|E_F T)$ (definitions of E_F , E_C)
7. $P_C(A|E_F) = C_0(A|E_F T)$ (6)
8. $P_F(A) = C_0(A|E_F T)$ (1, 6)

Alignment accordingly has the consequence that if one evidence base is a coarse-graining of another, then any agent which satisfies the PP for the chance function derived from the coarser evidence base *ipso facto* satisfies the PP for a chance function derived from the finer evidence base. Given Alignment, obeying a coarse-grained chance function and always updating by conditionalization on any finer-grained evidence *just is* obeying the finer-grained chance function obtained by conditionalizing the coarse-grained chance function on that finer-grained evidence.

¹¹ We first learned of the following argument from Cian Dorr. The formulation we give here is due to Chris Meacham.

Chris Meacham (this volume) considers the possibility of conflicts between the demands of different chance functions. Such conflicts would arise if two chance functions were to place jointly unsatisfiable constraints on rational initial credence. Then no initial credence function could count as rational, a disastrous result. The argument just given shows that Alignment prevents conflicts from arising between chances with evidence bases one of which supervenes on the other. We can also show, by generalizing an argument which goes back to (Lewis, 1980), that any theory which avoids such conflicts must incorporate Alignment.

If conflicts are to be avoided between chances at different levels, then the PPs for those chance functions must be jointly satisfiable. That is, there must be some rational initial credence function C_0 such that it delivers $P_F(E_FT)$ when conditionalized on E_FT , and delivers $P_C(E_CT)$ when conditionalized on E_CT . So we assume that PPs hold for both the coarser and finer chance functions, and consider the proposition X such that $(X \& E_C = E_F)$ ¹²:

1. $P_C(A) = C_0(A|E_CT)$ (PP_C)
2. $P_F(A) = C_0(A|E_FT)$ (PP_F)
3. $P_F(A) = C_0(A|E_CXT)$ (2, definition of X)
4. $P_F(A) = P_C(A|X)$ (1, 3)
5. $P_C(A) = P_C(A|E_C)$ (1)
6. $P_F(A) = P_C(A|E_CX)$ (4, 5)
7. $P_F(A) = P_C(A|E_F)$ (6, definition of X)

Together, these two arguments demonstrate that Alignment is both sufficient and necessary to prevent conflicts arising between chance functions one of which is a fine-graining of the other.

¹² Think of X as the proposition specifying, of the many fine-grained states compatible with the actual coarse-grained state, which one is actual.

Is Alignment a substantive principle of our theory of chance? In a sense, the answer is yes: there is no a priori guarantee that the world affords a non-trivial¹³ chance-theory T , and *a fortiori* there is no a priori guarantee that it affords multiple such theories related as per Alignment. So the existence conjecture for multi-grained chance is logically stronger than the existence conjecture for single-grained chance. In another sense, the answer is no: in requiring Alignment we have not expanded the chance role in any way. The above derivation of Alignment shows that all of the work is being done by the applicability of PPs at both the finer and coarser levels. Consequently, nothing we have said so far contradicts Lewis' claim that the PP captures all we know about chance.

A potential loophole remains: because supervenience is only a partial order over evidence bases, our framework allows for two chance functions derived from evidence bases E_1 and E_2 , neither of which is a coarse-graining of the other. In such a case, there will be no proposition X such that $X \& E_1 = E_2$. For an example, consider the relation that a chance function P_M with an evidence base E_M consisting of river-levels in Melbourne stands in to a chance function P_Y with an evidence base E_Y consisting of river levels all along the Yarra¹⁴. The evidence bases of these two chance functions partially overlap. That is, neither subject-matter supervenes on the other (E_M includes *the levels of Melbourne rivers other than the Yarra*, while E_Y includes *the levels of the Yarra outside Melbourne*) but the subject-matters are not independent (*the levels of the Yarra in Melbourne* is included in both).

In such a case, Alignment does not apply directly. So how can we be assured that P_M and P_Y do not impose conflicting constraints on rational credence? We must make use of the fact that both E_M and E_Y are coarse-grainings of a common evidence base which generates chances of its own: in our example, E_{SM}

¹³ By 'non-trivial' we mean 'having a grounding argument logically weaker than the whole truth about the world, past, present and future'. The omniscient function is trivial; all the other chance functions discussed in section 6 are non-trivial (if they exist.)

¹⁴ Geographical note: part of the Yarra flows through the Great Dividing Range, and part of it flows through Melbourne. Various other waterways, including the Maribyrnong River, also flow through Melbourne.

will do the trick. Because the microstate of the universe fixes (*inter alia*) the levels of all rivers everywhere, it fixes E_M and E_Y . The constraints Alignment places on the relations between E_{SM} and the river-level evidence bases allows us to derive a consistency constraint on chances deriving from the river-level bases:

1. $P_M(A|E_{SM}) = P_{SM}(A)$ (Alignment)
2. $P_Y(A|E_{SM}) = P_{SM}(A)$ (Alignment)
3. $P_M(A|E_{SM}) = P_Y(A|E_{SM})$ (1, 2)

That is, two coarse-grained chance functions with overlapping evidence bases must coincide once they are conditionalized on any evidence base corresponding to a genuine chance function of which each of their own evidence bases is a coarse-graining. Since neither of the functions conflict with a function which underlies them both, they cannot conflict with each other. In this sense, coarse-grained chances as we conceive of them cannot be autonomous.

We have shown that Alignment, which is necessary for the coherence of multi-grained chance, also places a strong constraint on how coarse-grained chances with partially overlapping evidence bases relate to one another. But recall that Alignment is not an additional component of the chance role: it flows directly from the PP. So chances characterized functionally via the PP cannot be autonomous. While this may limit their usefulness in strongly anti-reductionist philosophies of science, it helps to secure the credibility of coarse-grained chance in a more reductionist setting.

One further, terminological, issue needs to be mentioned. As we have formulated the multi-grained chance picture, a coarse-grained chance function is complete, in that it ascribes a chance to every contingent proposition. (This is a consequence of the existence conjecture discussed in §5.) Thus, the 'statistical-mechanical chance function' $P_{SM}(A)$ will assign chances to specific Bohmian trajectories constituting the microstates, and the 'poker chance function' $P_{POK}(A)$ will assign chances to specific atomic configurations constituting the cards. This might seem worrying; shouldn't coarse-grained chance have a limited coverage, so that a chance function ranges only over propositions wholly about its 'target' subject-matter? Well, we can talk this way if we like: then 'coarse-grained chance

functions' are individuated by pairs of subject-matters: the target subject-matter D, and the evidence base E. Such restricted-domain chance functions certainly exist. However, they will not in general stand in straightforward Alignment-style relations to one another, which limits their theoretical usefulness. But ultimately, 'coarse-grained chance function' is a technical term; what matters is how we account for chance-attributions in ordinary language.

This completes our sketch of the formal properties of coarse-grained chance. In the next section we investigate which sorts of coarse-grained chances are relevant to the semantics of ordinary chance-attributions.

9. Which Are Our Chances?

The framework as we have set it out so far places no restrictions on which partitions can represent evidence bases, and can thereby generate chance functions. For any contingent subject matter whatsoever, we can hypothesize that the global chance-theory T entails conditionals connecting maximal specifications of that subject matter to chance assignments to propositions. The examples of partitions based on the colours of objects described above would generate 'colour chances': probabilities of arbitrary propositions conditional on various facts about the distribution of coloured objects. But many evidence bases are more gerrymandered still: consider, for example, the evidence base comprising information about the total number of atoms in the fusion of the moon and the Pacific Ocean, or the evidence base comprising information about the names of all dogs with exactly one thousand fleas. Do such evidence bases generate genuine chance functions?

We will remain agnostic as to whether the true global chance theory T entails conditionals connecting gerrymandered evidence bases to values of coarse-grained chances. For our purposes, it doesn't matter: these chance functions, even if they exist, are clearly of no use in accounting for ordinary language uses of the word 'chance'.

Being gerrymandered is not the only way in which an evidence base might fall short of generating a chance function which is a reference of ordinary chance-talk. Consider *parochial chances*: chances with evidence bases restricted to particular spatio-temporal regions. Lewisian chances are a special case of this, where the spatio-temporal region in question is the entire universe up to some time.

However, we can also easily characterize Australian chances (probabilities conditional on how things are, timelessly, in Australia) and 17th century chances (probabilities conditional on the complete state of the world throughout the 17th century.) But like gerrymandered chances, parochial chances don't seem to be relevant to English chance-ascriptions¹⁵. Why not? Presumably, because the spatio-temporal symmetries of the world tend to render such chances less predictively useful than chances which are not spatio-temporally restricted in this way. A second consideration is that we rarely talk with others who share our parochial evidence bases.¹⁶ In the remainder of this section we attempt to make more precise the way in which context selects a particular subject matter as constraining the evidence base.

(Glynn, 2010) proposes that context supplies a 'level parameter'. We impose a similar requirement: for each context in which there are clearly-true chance-ascriptions, there is a contextually salient subject-matter which comprises an evidence base for a contextually salient chance function. Does this require positing an 'unarticulated constituent' in chance ascriptions? – that is, a meant but unspoken 'at level L' qualification? Not necessarily. We propose to import plasticity of fineness of grain from elsewhere in the semantics, by linking the variability of 'has chance x' to the variability of the modal verb 'could'.

We suggest that the salient evidence base at a context consists of the conjunction of propositions expressing the evidence which counts as *available* in that context, and that a type of evidence which counts as available is a type of evidence that *could be got by a particular method*. Our account can be factored into the following two components:

¹⁵ With the possible exception of Lewisian chance (discussed further in the next section).

¹⁶ Similarly, our proposal can model 'chances' generated from an evidence base consisting of propositions about mental states. Taking the evidence base to consist in an agent's total knowledge delivers Williamsonian *evidential probabilities* (Williamson, 2000). It is reflective of the degree of flexibility of our proposal that it can characterize evidential probabilities; but we suspect that these probabilities are of no special use in accounting for chance-ascriptions in English. We will not pursue the topic further here.

Contextualism: Context C fixes a chance function to feature in the semantics of chance-ascriptions by fixing the *available* subject matter S. Maximal propositions E, wholly about S, comprise the evidence base for the chance function salient in C.

Availability: 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.^{17, 18}

Although our usage of 'available' has much in common with the ordinary use of the word, we require that it always take a broad reading. We might ordinarily emphasize such a reading by saying 'in principle available'. But the details are unimportant: 'available' is acting as a semi-technical term for us, and what we take to be semantically basic is the connection with the modal 'could'.

When we speak with others, context determines that there are certain epistemic methods that are understood to be legitimate, possible, in use, or otherwise salient. Take the *maximal* amount of information that can be obtained by a given epistemic method M. This information will be equivalent to the cell of a partition: that partition is the subject matter 'what could be known by method M'. When playing poker, the salient method¹⁹ is one which does not involve violation of the rules of poker. Any salient method will be one which enables us to know the identity of all the visible cards. No salient method will enable us to know the microstate of the room in which the game is being played. Thus, what could be

¹⁷ We have formulated this principle in terms of what could be known (ignoring any Fitch-style complications), but we aim to remain neutral on questions about the metaphysics of epistemic states. For those who prefer not to build knowledge into metaphysical analyses (perhaps for internalist reasons), Availability could be reformulated in terms of information that could be acquired or evidence that could be obtained. If a non-factive notion is used, then the requirement that E be true becomes important.

¹⁸ Given the coarse-grained conception of propositions we are working with, it will turn out that necessary truths always count as available.

¹⁹ Or methods; we need not insist on uniqueness.

known in this context will be precisely the sort of knowledge which generates non-trivial poker chances.

What of the alleged ‘colour chances’? For there to be a conversation about colour chances would seemingly require a context in which the maximal information derivable from *all salient methods* tells us only about the distribution of colours, without telling us about anything else. Perhaps, if we invented a novel game in which players needed to keep track only of colour information, we could artificially induce a context in which such methods are salient. But putting those cases aside, we think that there is no reason to fear that we are inadvertently referring to colour chances in normal contexts. There are no reasonable epistemic methods that are contextually salient enough to render the colour chances eligible to be referred to.

We are assuming that modal expressions such as ‘could’ have contextually variable truth conditions. This has been fairly uncontroversial among philosophers following the work of Lewis (1976) and Kratzer (1977). Most familiar to philosophers, perhaps, is Lewis’s paper, which looks at variability of ‘can’, in examples such as: ‘David can speak Finnish, because he has the right sort of larynx’ and ‘David cannot speak Finnish, because he has never learned the language’. The central thought is that to say something *can* happen is to say that it is compossible with certain facts, but that the relevant facts may vary with different contexts.

There is a notable precedent for our appeal to the context-sensitivity of ‘could’ in characterizing chance. (Eagle, 2011) argues for the following principle linking chance and the closely-related word ‘can’:

Chance-Ability Principle: Where X is a noun phrase and φ a complement verb phrase, the chance of X φ -ing exceeds zero iff X can φ .

According to our proposal, α has some chance of occurring iff the evidence that could be got about α does not rule out that α occurs. According to Eagle’s proposal, α has some chance of occurring iff α can occur. Are these proposals independent of one another? It depends on whether ‘ α can occur’ and ‘the evidence that could be got about α doesn’t rule out that it occurs’ are independent of one another. While

the latter doesn't entail the former (perhaps there is unavailable evidence which rules out α happening) it does seem to us that the former entails the latter: *if α can happen, then no evidence could be got which rules out α happening.*

If this is correct, then all cases which Eagle's proposal rules as non-trivially chancy will also count as non-trivially chancy according to our proposal. That is, the constraint which Eagle's proposal places on the semantics of chance-attributions is guaranteed to be met if our proposal holds. The evidence that Eagle adduces in favour of CAP is thus successfully explained by our proposal.

A feature of our account is that while 'there is some chance of α , though we can't know what it is' is perfectly consistent, 'there is some chance of α , though we couldn't know what it is' is inconsistent.²⁰ While these consequences might perhaps be used as the basis of an objection, we choose to embrace them. The distinction between 'can' and 'could' does seem to be recognized by ordinary speakers. There are some things that we could do but in fact can't. Consider the following everyday exchange: "Can you do it?" - "No; I could, but I don't have time."

The proposal set out in this section entails that, whatever the context, chance-attributions to contingent propositions are always meaningful. For example, even in a context where the poker evidence-base is salient, it would make sense to ask what the chances are that World War Three will break out by 2050. The act of asking such a question might well have the pragmatic effect of changing the context so that some geopolitical evidence base is salient; but the chance is well-defined (even if unknowable) whatever the context. This 'universalism' marks a point of departure from Hoefer's account; Hoefer maintains that ordinary chance-attributions refer to chance functions with highly-limited domains, and consequently rejects Alignment in favour of a patchwork-style view of chance-

²⁰ At least, it comes out inconsistent unless we allow for in-sentence context-shifting. Perhaps there are contexts where such context-shifting is appropriate. For example, we may want to allow a true reading of 'I won't win the lottery, though I suppose it's possible'. However, such context-shifting is generally seen as a last resort in philosophical semantics; we want to emphasize that our account does not rely on it to account for any core cases.

theories. Lewis was apparently cautious (Lewis, 1980, pp.276-7) about allowing all contingent propositions to be assigned chances; however, as (Meacham, 2007) emphasizes, Lewis' concern here related specifically to problems about infinite state-spaces and chance-zero events. Lewis' account is universalist in spirit, as he makes clear in (Lewis, 1986, p.132); we follow him in this, setting aside the problem of chance-zero events.

10. Against Lewisian chance

It is obvious why the *Lewisian* chance function should have been of interest to philosophers. In discussions of free will and determinism metaphysicians have frequently had cause to consider whether the entire current description of reality suffices to determine the future. But such discussions take place in specialized contexts, and we should be wary of over-hasty generalization from them. In this section, we will discuss which chance functions are best placed to feature in the truth-conditions of our ordinary chance-talk.

(Schaffer, 2007) argues that the Lewisian (fine-grained time-varying) chances best satisfy various platitudes about chance. Before responding directly to these arguments, we offer a compatibilist platitude of our own: gambling games will remain paradigm games of chance, however fundamental physics should turn out. (Eagle, 2011) offers a further platitude that supports the compatibilist approach: where A can do B, there is some chance of A's doing B. Any account of chance which aims to do justice to the folk conception of chance should capture these platitudes, or at the very least should explain them away.

The first of Schaffer's arguments turns on the Principal Principle. Schaffer's discussion presupposes that all historical information is admissible, and he regards the claim that some historical information is inadmissible as a *revision* of the PP. If we were to identify the PP with L2PP, this description would make sense; but Schaffer takes the PP to be a platitude about chance rather than a theory-laden principle. And it is highly implausible that it is a platitude that historical information of arbitrary detail is admissible in all contexts. Here, as elsewhere, Schaffer motivates a principle by a plausible informal argument, but then goes on to give a precise formulation of the principle which incorporates additional

unargued-for commitments. This will be a recurring theme in our discussion of his platitudes.

As well as making the unjustified charge of revisionism against compatibilist treatments of the PP, Schaffer extends a challenge. The compatibilist:

... will need to explain why her revised principle bears any real connection to objective chance, especially when setting one's credences to it will make one endorse bets one knows are doomed. (Schaffer, 2007) p.129

This objection is easily met in the current framework. It is indeed irrational to have credences which diverge from the fine-grained chances when those chances are known: but this does not mean that the compatibilist who knows the fine-grained chances must violate the coarse-grained PP. The argument given at the end of §7 shows that, if one knows the fine-grained chances, one will also know information which, when conditionalized upon, results in the convergence of the advice given by coarse-grained chances and fine-grained chances.

Schaffer's second platitude is the Realization Principle (RP), which requires that, if some outcome has a non-zero chance, then it is a genuine possibility. Stated informally, there is no reason to think that compatibilist chance cannot meet this criterion. As with the PP, it is in Schaffer's formal version of the principle that incompatibilism is smuggled in:

Realization Principle (RP): If $ch\langle p, w, t \rangle > 0$, then there exists a world w_{ground} such that: (i) p is true at w_{ground} , (ii) w_{ground} matches w in occurrent history up to t , (iii) w_{ground} matches w in laws. (Schaffer, 2007) p. 124

It is a small matter to replace 'occurrent history up to t ' with 'occurrent history with respect to evidence base E up to t ' in this formulation. The resulting formulation enjoys at least as much intuitive support as Schaffer's formulation (that is: not much!) and it is satisfied straightforwardly by compatibilist chance.

Schaffer's third platitude is the Future Principle, which says that non-trivial chances pertain only to the future: in other words, that past events are no longer chancy. It is easy to specify versions of compatibilist chance which satisfy the

Future Principle, but it is not clear that we ought to impose this requirement in the first place. Even leaving aside exotic time-travel cases, ordinary discourse seem to allow for non-trivial chance ascriptions to events known to be in the past. Consider the following example:

Footy match: Before a match I judged that the chance of a St. Kilda win was $\frac{1}{2}$. Now the match has finished, but I have not yet heard the score.

An appropriately coarse-grained PP constrains my before-match credence in a St. Kilda win to be $\frac{1}{2}$. Clearly I should retain credence $\frac{1}{2}$ in a St Kilda win after some time has passed and no new evidence has come in, but this can be explained by both the timeless and time-varying accounts. We either apply the PP pre-match and then retain credence $\frac{1}{2}$ because we have no new relevant information; or we apply the PP directly post-match. The question which distinguishes these different accounts is whether, after the match, I can still correctly judge that the chance of a victory is $\frac{1}{2}$. We suspect that ordinary speakers are inclined to think that ascriptions of chanciness to the past are acceptable in many contexts, but we shall not attempt to argue for this conclusion here. (It seems in any case a matter more for empirical investigation than for philosophical intuition-eliciting.) We note only that our generalized framework for chance can accommodate both possibilities, while Schaffer's approach cannot.

Schaffer's fourth platitude, which he calls the Intrinsicness Requirement (IR), is that chances should remain the same across chance setups that are intrinsic duplicates of one another. The compatibilist version of this principle extends intrinsic duplication from exact intrinsic duplication to intrinsic duplication with respect to a particular subject-matter. Compatibilist chance straightforwardly satisfies the IR thus extended. We suggest that our formulation enjoys no less intuitive support than Schaffer's formulation.

Schaffer calls his fifth platitude the Lawful Magnitude Principle: informally, it amounts to the requirement that chances should be determined by laws of nature. Our compatibilist picture is fully compatible with this principle; indeed, it in fact accommodates it better than Schaffer's incompatibilist picture. Probabilities determined by fundamental laws of nature count as chances according

to both compatibilism and incompatibilism; however, the incompatibilist cannot also allow that the probabilities derived from statistical mechanics are genuine chances, while the compatibilist can. Presumably Schaffer will claim that statistical mechanics does not involve any genuine laws of nature. In our view, this does not accord well with scientific usage: physicists seem to count statistical-mechanics as specifying genuine, and genuinely probabilistic, laws.

The sixth and final platitude is the Causal Transition Constraint. Schaffer glosses this as the requirement that ‘chances should live within the causal transitions they influence’. As (Glynn, 2010) notes, Schaffer’s formulation of this constraint is subject to apparent counterexamples; and anyway, it presents no obvious challenge to compatibilist chance. We will set it aside.

The upshot of our discussion is that none of the platitudes themselves (as opposed to Schaffer’s tendentious formulations of them) present any problem for compatibilism. But a stronger conclusion is possible. As well as merely demonstrating compatibility with the platitudes, we would like to explain why the platitudes hold. Incompatibilists face a particular problem here with respect to the Principal Principle: when chances are limited to primitive propensities at the fundamental level, the project of connecting them to rational belief looks utterly intractable. By contrast, our compatibilist picture offers a promising approach to justifying the PP. It is not surprising that language-users would develop a modal idiom which reflects the limits of what can be inferred from available bodies of information, even without knowledge of how that information relates to goings-on at the fundamental level. In (Handfield, MS), one of us has proposed an explicit genealogy of our chance-talk along such lines.

Our proposal brings out the commonalities between objective probabilities from different domains, in a way which the incompatibilist orthodoxy does not. What different chance functions have in common is that they are generated by counting as inadmissible domains of facts which couldn’t be known but which nonetheless are relevant to facts which could be known. Where the chance functions differ is with respect to the reasons why the inadmissible facts which would trivialize them could not be known. The poker-chance-trivializing facts couldn’t be known without violating the rules of the game; the SM-chance-

trivializing facts couldn't be known without our being thoroughly different (non-thermodynamic) creatures; the Bohmian-chance-trivializing facts couldn't be known without our being even more thoroughly different (sub-quantum) creatures. This unified account of chance contrasts well with the incompatibilist treatment, with its sharp dualism between primitive fundamental chances and derivative merely-epistemic probabilities.

11. Conclusions

All of the elements of our account are now in place. Recall the principle, characteristic of incompatibilism about chance, with which we started:

Chance Exclusion Principle: If 'the chance of ϕ is x ' is true at t , then 'the chance of ϕ is y ' is true at t iff $x=y$.

On our view, CEP is false, although it is true when restricted to chance-attributions uttered in the same context as one another. It fails when applied to chance-attributions uttered in contexts over which the salient evidence base is allowed to vary. One cannot legitimately infer a contradiction from the honest gambler truly saying to a friend that his chance of winning is one half, and the cheating gambler truly saying to a friend that the chance of the honest gambler winning is zero.

This feature of our proposal allows for a powerful response to a prominent incompatibilist line of argument. Consider the following argument:

1. The chance that the fair coin landed Heads is $\frac{1}{2}$. (Assumed for *reductio*)
2. The coin landed either Heads or Tails. (Premise)
3. If the coin landed Heads, then the chance that it landed Tails is zero. (Premise)
4. If the coin landed Tails, then the chance that it landed Heads is zero. (Premise)
5. The chance that the coin landed Heads is either zero or one. (From 2, 3, 4)
6. *Reductio*. (From 1, 5)

Our diagnosis is that this reasoning exploits a context shift. Premises 1 and 5 are never true in the same context. According to our proposal, premise 1 is true in just those contexts where information about the result of the coin is unavailable. But entertaining premises 3 and 4 moves us out of those contexts. The antecedents of 3 and 4 are explicit suppositions about the way the coin landed; making such suppositions tends to move us into contexts where information about the way the coin landed is (in principle) available.

As Lewis famously remarked, ‘it is no fault of a context that we can move out of it.’ (Lewis, 1979). However, as in other applications of contextualism (contextualism about knowledge is the most obvious example), it seems to be easier to change contexts in one direction than the other. Once fine-grained chances are salient, it can be hard to make coarse-grained chances salient again. This phenomenon can be explained, on our proposal, in terms of the greater informativeness of fine-grained chances.

We have shown how the generalized framework for chance described by Arntzenius, Meacham, and others reconciles determinism with non-trivial chances. Coarse-grainings of the evidence base generate a partial order of chance functions, with finer trumping coarser if known. We have proposed that context determines which chance functions are relevant for assessing particular chance-attributions. And we have suggested a mechanism by which this contextual determination of chance function could operate, a mechanism which appeals only to the comparatively familiar context-sensitivity of ‘could’. The incompatibilist conception of chance which dominates in contemporary metaphysics can be seen as a special case of our generalized conception. We believe that the resulting framework for chance does justice to our variegated use of chance-talk in everyday life, in the special sciences, and in fundamental physics.

References

- Armstrong, F. (1995). *Chance and the Principal Principle: Things Ain't What They Used To Be*.
- Armstrong, F., & Hall, N. (2003). On What We Know about Chance. *British Journal for the Philosophy of Science*, 54(2), 171-179.
- Bigelow, J., Collins, J., & Pargetter, R. (1993). The Big Bad Bug: What Are the Humean's Chances? *British Journal for Philosophy of Science*, 44(3), 443-462.
- Briggs, R. (2009). The Anatomy of the Big Bad Bug. *Noûs*, 43(3), 428-449.
- Eagle, A. (2011). Deterministic Chance. *Noûs*, 45(2), 269-299.
- Glynn, L. (2010). Deterministic Chance. *British Journal for the Philosophy of Science*, 61(1), 61-80.
- Hájek, A. (2003). What Conditional Probability Could Not Be. *Synthese*, 137(3), 273-323.
- Hájek, A. (2007). The Reference Class Problem is Your Problem Too. *Synthese*, 156, 185-215.
- Hall, N. (1994). Correcting the Guide to Objective Chance. *Mind*, 103(412), 505-517.
- Handfield, T. (MS). *Chance: The Non-Existent Best Advice*.
- Hofer, C. (2007). The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance. *Mind*, 116(463), 549-596.
- Ismael, J. (2008). Raid! Dissolving the Big, Bad Bug. *Noûs*, 42(2), 292-307.
- Lewis, D. K. (1979). Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8(1), 339-359.
- Lewis, D. K. (1980). A Subjectivist's Guide to Objective Chance *Studies in Inductive Logic and Probability*, V2.
- Lewis, D. K. (1986). Postscripts to 'A Subjectivist's Guide to Objective Chance'. In D. K. Lewis (Ed.), *Philosophical Papers Vol. II* (Vol. I, pp. 114-132). Oxford: Oxford University Press.
- Lewis, D. K. (1988a). Relevant Implication. *Theoria*, 54(3), 161-174.
- Lewis, D. K. (1988b). Statements Partly About Observation. *Philosophical Papers*, 17(1), 1-31.
- Lewis, D. K. (1994). Chance and Credence: Humean Supervenience Debugged. *Mind*, 103(412), 473-490.
- Loewer, B. (2001). Determinism and Chance. *Studies in History and Philosophy of Modern Physics*, 32, 609-620.
- Meacham, C. (2005). Three Proposals Regarding the Theory of Chance. *Philosophical Perspectives*, 19(1), 281-307.
- Meacham, C. (2010). Two Mistakes Regarding the Principal Principle. *British Journal for Philosophy of Science*, 61(2), 407-431.
- Popper, K. (1956). The propensity interpretation of probability. *British Journal for the Philosophy of Science*, 10, 25-42.
- Schaffer, J. (2003). Principled Chances. *British Journal for the Philosophy of Science*, 54(1), 27-41.
- Schaffer, J. (2007). Deterministic Chance? *British Journal for the Philosophy of Science*, 58(2), 113-140.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.