

*Raid! Dissolving the Big, Bad Bug*¹

JENANN ISMAEL

Centre for Time, University of Sydney,
Department of Philosophy, University of Arizona**Introduction**

There's a long history of discussion of probability in philosophy, but objective chance separated itself off and came into its own as a topic with the advent of a physical theory—quantum mechanics—in which chances play a central, and apparently ineliminable, role. In 1980 David Lewis wrote a paper pointing out that a very broad class of accounts of the nature of chance apparently lead to a contradiction when combined with a principle that expresses the role of chance in guiding belief. There is still no settled agreement on the proper response to the Lewis problem. At the time he wrote the article, Lewis despaired of a solution, but, although he never achieved one that satisfied him completely, by 1994, due to work primarily by Thau and Hall, he had come to think the problem could be disarmed if we fudged a little on the meaning of 'chance'. I'll say more about this below. What I'm going to suggest, however, is that the qualification is unnecessary. The problem depends on an assumption that should be rejected, *viz.*, that using information about chance to guide credence requires one to conditionalize on the theory of chance that one is using. I'm going to propose a general recipe for using information about chance to guide belief that does not require conditionalization on a theory of chance at any stage. Lewis' problem doesn't arise in this setting.

What's at stake?

It's easy for the philosophical issues to be buried in technical exchanges or debates over principles that are barely intelligible to the unaffiliated, but metaphysical issues of quite central importance hinge on resolution of Lewis' problem. The central difficulty can be put quite plainly: it looks like nothing that plays the role of chance in guiding belief could itself either *be*, or supervene on, local matters of particular fact. That is a very powerful conclusion, and the argument for it looks as tight as it can be. Lewis formulated a principle that was intended to capture the role in guiding belief, combined it with the hypothesis that chances supervene on local matters of particular fact (really, a stronger premise that we'll see below), and derived a contradiction. If the argument is correct, the whole analytic project that Lewis had been carrying out piecemeal by offering accounts of concepts that were thought to be problematic for the defender of the principle that everything supervenes on local matters of fact would be dashed. He had already given largely successful accounts of laws, causes, and counterfactuals, in each case employing a reductive strategy tendering patterns in the global manifold of events as truthmakers for the problematic structures. Lewis thought that if you apply the same treatment to chance, you get trouble.²

Taking over Lewis' stage setting, one assumes that the world consists of events related by dynamical laws that give the general bearing of the past on the future. These dynamical laws find expression in history-to-chance conditionals, which generate at any point in time, a function that assigns a probability to any event on the basis of history up to that time.³ For any time t and metaphysically possible event e , we can ask "what is the probability of e at t ?"⁴ If e is in t 's past, the answer will be either 0 or 1. If it lies in t 's future, the answer may be any real value in the interval between 0 and 1.

The Principal Principle (hereafter, PP) is a principle intended to capture the role that chances play guiding belief. Lewis believed that it captured all we know about chance,⁵ and expressed it thus in prose:

"If a rational believer knew that a chance of [an event e] was 50%, then almost no matter what he might or might not know as well, he would believe to degree 50% that [e] was going to occur. Almost no matter, because if he had reliable news from the future about whether e would occur, then of course that news would legitimately affect his credence."⁶

PP says that if you know what the chance of an event is, and you don't have magical access to the future, you should adopt it as your credence, and it will screen off (render probabilistically irrelevant) all other information, save reliable news from the future. The lack of an explicit, non-circular definition of "reliable news from the future" has been a source of concern since the principle was proposed. If the only way we have of characterizing reliable

news from the future is as ‘information that overrides chance’, the principle is empty. It says simply ‘adjust credence to chance whenever you don’t have information that overrides chance’. We can resolve this difficulty by replacing ‘reliable news from the future’ with ‘information about the future that cannot be obtained from natural sources, and helping ourselves to a physical understanding of the basis of natural sources of information about the future. This is a somewhat complicated story, but we say that the only information that overrides the chances is information that can’t be derived from the events in one’s past light cone by application of dynamical laws, assuming a canonical statistical distribution over initial conditions.’⁷

In his original presentation, Lewis assumed that history and the theory of chance were both admissible and formulated PP;

$$\mathbf{PP}_{\text{orig}} \text{Cr}_t(A/H_t T_w) = \text{Ch}_t(A),$$

The subscript stands for ‘original’. ‘Cr’ denotes credence, ‘Ch’ denotes chance, ‘H’ refers to historical information and ‘Tw’⁸ refers to the set of history-to-chance conditionals, hereafter, the ‘theory of chance’. This should be read “one’s credence in A at t, given pre-t history and the theory of chance, is the chance of A at t”.

Suppressing the temporal subscripts

$$\mathbf{PP}_{\text{orig}} \text{Cr}(A/HT_w) = \text{Ch}(A)$$

Now let’s suppose that the truth conditions for statements about chances defined at a point are distributed patterns defined over the whole manifold, and that the chances defined at a time don’t supervene just on past history.⁹ We’ll call this the Humean Hypothesis.

The argument that the Principal Principle conflicts with the Humean Hypothesis is simple. If the Humean Hypothesis is true, so long as there are any chancy events, there will be futures that have a non-zero chance of occurrence but which conflict with the present chances in the strong sense that they literally contradict the claims that we have tacitly made about the future in describing the present chances. Lewis put this by saying that there are future courses of history that are allowed by the physical laws which *undermine* the present chances. Let F be such a future. By $\mathbf{PP}_{\text{orig}}$, $\text{Ch}(F/E) \neq 0$, but A is *inconsistent* with E, so $\text{Ch}(F/E) = 0$. Contradiction.

We cannot apply $\mathbf{PP}_{\text{orig}}$ to determine our subjective probability for undermining futures without assigning a non-zero probability to contradictions. As Lewis put it:

“our problem, where F is an unactualized future that would undermine the actual chances given by E is that $\text{Ch}(F/E) = 0$, because F and E are inconsistent, but

Ch(F/E)≠0 by **PP**_{orig} because E specifies that F has some present non-zero chance of coming about.”¹⁰

It was clear enough to Lewis from the outset what was going on. Any plausible Humean account, i.e. any that satisfies the Humean Hypothesis, is going to entail that the present distribution of chances contains information about the future, specifically, the information that undermining futures won't occur. This is because part of the global pattern in events in virtue of which T_w obtains, lies in the future, and this means that the present distribution of chances, derived from T_w in combination with history, contains such information. If the Humean Hypothesis is correct, it turns out that the theory of chance itself is a kind of crystal ball. On the one hand, we know that the presence of crystal balls defeated applications of **PP**_{orig}, but on the other hand, as Lewis argued, we can't block the case by ruling T_w inadmissible, for if T_w is inadmissible, then *no* application of **PP**_{orig} is legitimate. We might as well say “chance is guided by credence except on days on which the sun rises, the tides come in, or somebody eats a french fry”. Lewis regarded this neither as a challenge to the consistency of **PP** nor to its claim to correctly capture the concept of chance. He regarded it as a *reductio* of any plausible Humean account of the nature of chance. He regarded it as demonstrating that nothing in a Humean world can play the role precisely defined by **PP** in guiding belief. Bad news not for chance, but for Humean metaphysics. Looking back fourteen years later, Lewis wrote of the argument:

“If I'd seen more clearly, I could have put the core of my *reductio* thus. . . information about present chances is inadmissible because it reveals future history. But this information is not inadmissible, as witness the way it figures in everyday reasoning. Contradiction.”¹¹

His own remedy was to accept an emended principle.¹² The result, when formulated in a general way, is a more complex principle that doesn't precisely capture the pre-theoretical conception of chance, which Lewis continues to maintain is captured by **PP**_{orig}, but one that is both generally applicable and consistent, and to which the old principle is a good approximation except under extraordinary conditions.¹³ The way to assess the situation philosophically is that there is nothing that plays the role of chance in a Humean world, but there is something closely enough related to explain how we got the concept of chance, and this closely related quantity is what guides belief. So for Lewis, there is a heavy concession in accepting the emended principle.

To recap the logic:

Humean Hypothesis • for any time t , so long as there are chancy events at t , there exist future courses of history that have a positive probability of occurrence but that are metaphysically incompatible with the chance distribution at t .

Humean Hypothesis + PP → contradiction.

Humean Hypothesis + NP → no contradiction.

Lewis accepts demonstration of consistency of Humean Hypothesis with NP as a resolution, but only by conceding to the non-Humean that there are no chances, but only a close proxy, in a Humean world.

The New Proposal

This is a familiar story, recounted by Lewis and rehearsed by others.¹⁴ I submit, against near perfect orthodoxy, that rather than making the principle more complex to avoid inconsistency, we should simplify it. It was a mistake to conditionalize on the theory of chance in the first place. Lewis' verbal gloss captured the concept quite nicely and it didn't mention theories of chance, or history to chance conditionals, or anything of the sort. The chances are those objective features of the world, whatever they may be, to which one ought to adjust one's subjective probabilities, no matter what historical information she possesses.¹⁵ They are those objective features of the world, that is to say, knowledge of which screens off knowledge of all information from the past. And that in its turn just says that the chance of A at p should determine credence in A at p, conditional upon information about the contents of p's past light cone:

$$\mathbf{PP}_{\text{uncond}} \text{Cr}_p(A/H_p) =_{\text{def}} \text{Ch}_p(A)$$

As an implicit definition of the concept of chance, this does much better, it doesn't have problems with undermining, and its epistemic motivation is plain. In a world in which information is spread around in the way it is in ours, the only conditions under which a situated observer would be wise *not* to adjust credence to chance is if she has information *not about*, but *from* the future, information of a kind that *we* do not come by through natural means.

Reasoning within the scope of a theory vs. reasoning about theories

Why, then, did Lewis, and nearly everybody that followed him, conditionalize on the theory of chance? There is a confusion, very easy to make, between reasoning within the scope of a theory and reasoning *in propria persona* about which theory is correct. A contradiction can be obtained from the Humean Hypothesis only if one simultaneously uses T_w to assess the probability of an undermining future and conditionalizes on the truth of T_w . That never happens; it doesn't happen when one is reasoning within the scope of T_w because there is no need under those conditions to conditionalize on its truth. But to assign a positive probability to an undermining future is to assign a

positive probability to something that entails its own falsity. Isn't a theory that assigns a positive probability to something that entails its own falsity inconsistent? No, and here is where the temporal subscripts are important. A theory of chance assigns a positive chance to undermining futures conditional on any finite segment of history, but that is perfectly consistent with maintaining that T_w is not only true but *necessarily* true, given total history. And one doesn't conditionalize on a theory of chance when reasoning *in propria persona* because one is always, under those conditions, dividing belief between theories.¹⁶

Let me spell this out; when reasoning within the scope of a theory of chance, the theory is embodied in one's practice rather than included in the basis for conditionalization. To reason within the scope of T_w is to use it to govern one's expectations for the future; it is to adopt it as a policy.¹⁷ One's calculations are predicated on its truth, but not in a sense that requires its inclusion in the basis for conditionalization. Compare: an inference from p and $(p \rightarrow q)$ to q is predicated on the truth of $(p \ \& \ (p \rightarrow q) \rightarrow q)$ but not in a sense that requires one to take it as a premise of one's reasoning. The difference between epistemic policies and conditional beliefs is dramatized nicely in the famous exchange between Achilles and the tortoise in which Achilles tries to force the tortoise's assent to the conclusion of an instance of *modus ponens*. The futility of the exercise makes it clear that no belief or set of beliefs of the form $A \rightarrow B$ will substitute for the policy of inferring B from A , and conversely the policy of inferring B from A doesn't require the corresponding conditional belief.¹⁸

Once reasoning within the scope of a theory of chance is distinguished from reasoning about which theory of chance is *true*, we apply familiar principles for accommodating theoretical uncertainty by dividing opinion between the results obtained from different theories. And things here turn out just as they should. So long as our evidence is always drawn from history, we are never in the position of conditionalizing on the truth of a theory of chance, and we just don't get a conflict.

There are two principles at work in forming expectations in the face of uncertainty, there is the real, unconditional form of the Principal Principle:

PP_{uncond}: always and everywhere adjust credence to chance, no matter what historical information you possess; $Cr(A/H) = Ch(A)$.

And then there is the general principle for accommodating ignorance about the correct theory of chance:

Ignorance Principle: where you're not sure about the chances, form a mixture of the chances assigned by different theories of chance with weights determined by your relative confidence in those theories.

Every agent, in a world like ours, not possessed of non-natural sources of information, is using information from the past to do two things simultaneously: (i) to assess the relative probabilities of epistemically possible theories of chance, and (ii) to determine the chances assigned by the various theories. A theory of chance is a rule for transforming information about the past into a chance distribution.

PP_{uncond} is the key to our concept of chance; the ignorance principle applies general principles for reasoning in the face of uncertainty. Combining the two, we get a recipe for calculating credence in the face of uncertainty about the chances:

$$\text{The General Recipe : Cr}(A) = \sum a_{T_w} \text{Ch}_{T_w}(A).$$

where $\text{Ch}_{T_w}(A)$ is the chance of A according to T_w and a is your subjective assessment of the probability of T_w . The general recipe says, effectively, that you should adjust credence in A to your best estimate of the chances.¹⁹

The ignorance of someone employing the General Recipe about which theory of chance is true will be reflected quite precisely in the credences she assigns. The a_{T_w} 's here represent credence in the associated theory of chance and they should reflect the likelihood that the theory has, given the evidence at the agent's disposal, which is to say, how well it fits the frequencies. The General Recipe is only as good as an agent's procedures for determining the values of the a_{T_w} 's. If those procedures are good ones, we can expect credences assigned by the General Recipe to converge towards the chances. If not, not. Since reasoning about which theory of chance is true is reasoning that, by its nature, occurs outside the scope of any theory of chance, chances won't be a guide here. I openly acknowledge that the General Recipe is just a schema about how to use chance to guide credence under uncertainty until supplemented with an account of how to weigh evidence against different hypotheses about what the chances are. Giving such an account raises challenging issues about sampling techniques and reference classes that demand specialized treatment, and that can be bracketed for purposes of establishing the consistency of the Humean Hypothesis and the Principal Principle. For those purposes, we need only the weak prescription that an agent should recognize as epistemic possibilities all theories of chance that are logically compatible with the evidence. She should assign a weight of 0 only to those theories that are definitively ruled out by observed frequencies (which is to say, in practice, to *no* theory, for the only way a theory can be logically incompatible with observed frequencies is by assigning a probability of 0 (or 1) to an event that is observed (or not), but even in such a case, there is always an ineliminable *possibility* of error).

What of Lewis' worry? Information about the present chances *must* be admissible otherwise how could we adjust our credence about future events to the present chances? Doesn't this require that the facts about the world

in virtue of which the present chances are what they are, are available to agents who have epistemological access only to the past, and hence that no plausible Humean analysis of chance can be correct? No, we now see an alternative; allow that the chance-making pattern at any point in time lies partly in the future, and observe that we are precisely as ignorant of the present chances as we are of the remainder of the chance-making pattern (specifically, of which undermining futures won't occur). This doesn't keep us from using our best current estimate of the chances to guide expectation. Ignorance about the future (specifically, about which futures would, were they to occur, undermine the present chances, from a Humean point of view) is properly represented as uncertainty about the present chances. Compare: we are precisely as ignorant of the physical laws at any point in time as we are of which law-undermining futures won't occur. This doesn't stop us from using our best current estimate of the laws to derive predictions.

We identify chances with objective features of the physical manifold, explain why those features play the role they do in our own belief forming procedures by adversion to facts about our epistemological access to the future, tacit recognition of which is elicited by suspension of **PP**_{orig} in imagined cases in which we extend our epistemic powers with crystal balls and magical insight. The epistemology makes the real world extension of chance, though perfectly well defined, impossible to ascertain with certainty from any point within history, but it doesn't mean we can't have information about chances that falls short of certainty and it doesn't mean that information about chance can't guide belief. Compare: I know what a good stock pick is in the sense that I know the truth conditions for statements of the form 'x is a good stock pick', and I want my beliefs about good stock picks to guide my financial decisions. Those truth conditions lie partly in the future and that means that I am never certain until after the fact, whether a given pick is a good one. But it doesn't mean that I can't presently have information about good stock picks, and that that information—appropriately qualified—cannot guide decisions. How do I make financial choices? I consult the experts, paying attention to their track records, credentials, and the information they have access to. I watch market indicators and try to discern trends. There is a wealth of information contained in the past and present, whose informativeness is unimpugned by the fact that it falls short of certainty. A wise investor will make use of all of it. And a wise agent using information about chance to guide belief will do the same.

We can maintain direct and unqualified connection between chance and credence captured by **PP** in its unconditional form while maintaining a qualified commitment at any point *in time* to hypotheses about which theory of chance is true. The commitment is qualified because we are missing relevant information. The non-Humean cites incomplete knowledge of present and past as the source of the missing information. The Humean cites ignorance

of the future. But they accommodate their uncertainty in the same way. **PP** says “let chance guide credence”; this in its turn goes into “let T_w guide credence/ T_w is the correct theory of chance at our world (i.e., that T_w , together with history, gives the correct assignments of chance at our world)”, and that counsels us to let T_w guide credence only to the extent that we have confidence in T_w . Uncertainty enters not in the relation between chance and credence, but in the relation between T_w and chance.

So, two insights are needed for a resolution of Lewis’ worries. The first is that when one is using a theory, the theory shouldn’t be included in the basis for conditionalization. The second is that our ignorance of the future tempers our knowledge of the present chances in a way that quite precisely keeps us from certainty about the truth of any particular theory of chance. Every present degree of probability we assign a T_w -undermining future is one we take away from T_w , and hence the distribution of present chances derived from T_w . So, the conceptual relation between chance and credence, the one that in Lewis’ mind implicitly defined the concept, is given by the simple, unconditional prescription; always and everywhere adjust credence to chance when you know it²⁰ and combined with ordinary principles for mixing probabilities to yield a principle one can use to guide credence in the face of ignorance. No restriction of any sort is needed on the revised principle, **PP**_{uncond}, it can be consistently combined with a Humean Hypothesis and maintained without qualification.²¹

There is a nice illustrative scenario drawn from the paper by Hall containing the revised version of the Principle that Lewis himself accepted, which we can use to illustrate the General Procedure. Hall envisions a panel of putative experts, each of whose credence functions embody a theory of chance.²² If we ask a panelist at any time about some particular event, she will give us a probability. Assuming that at any given time, the panel contains all and only experts whose credences have not been definitely ruled out as chances, among the panelists, there is a real expert—‘Cassandra’—who embodies the correct theory of chance. In these terms, to reason within the scope of a theory of chance is to adopt as our own the credences recommended by the expert who embodies that theory. If we have no magical insight into the future, we won’t know who the real expert is after any finite bit of history, and we take account of this fact in our reasoning by forming a mixture of the opinions of different panelists, with weights determined by their track record, how well their credences have tended to fit the frequencies (acknowledging that there are unsettled questions about how this is gauged).

Modesty

One might yet worry that a theory that—as it is often put—assigns itself a non-zero probability of falsity is somehow inconsistent. The first thing to say is that theories of chance do not assign themselves probabilities. They

are not themselves events, and they simply fall outside the field over which chance is defined. This leads to confusion in the literature, in which the practice of flipping back and forth between the probability a theory assigns to an undermining future and the probability it assigns to itself is made carelessly, temporal subscripts are dropped, and talk of the probability that a theory of chance assigns to itself is customary. We could assume that a theory of chance that provides a complete distribution over the space of possible worlds, implicitly assigns itself a chance. But there are two reasons to resist this. First, the assumption that a theory of chance provides complete distributions over the space of possible worlds is less than innocent. Quantum mechanics, our central example of a theory of chance from physics, assigns probabilities to results of possible measurements, on the basis of the present state of a system and whatever external influences it is subject to,²³ but on the standard interpretation one cannot at any given time assign probabilities to events that occur in the non-immediate future without knowing the results of all intervening measurements. Even putting that aside, a theory of chance that assigns probabilities to finite strings of future events will typically be silent on total histories, not because history is necessarily infinite, but because for any finite string of events, there is always the *possibility* of events to follow. Short of an assumption that the universe has a history of a finite, specified duration, there is no general way of turning a distribution over finite strings of future events into a distribution over total histories. And without an assignment of probabilities to total histories, we don't have an assignment of probabilities to the Humean truthmakers of theories of chance.

We can illustrate this with the example of an actual frequentist account of chance. The actual frequentist holds that the chance that, e.g., a photon prepared in state ψ passes a polarizer at angle α is the overall actual frequency in w with which photons in ψ pass polarizers set at α . T_w can assign a positive probability to every single photon in ψ passing every single photon at α for any finite initial segment of w 's history without assigning a probability to an overall frequency, because for every finite string of events, there is the possibility of events to follow in which frequencies diverge as wildly as you please. To get an overall frequency, we need to add that *there are no other events*, and that is not something that chances generated by history-to-chance conditionals will in general, assign a probability. The same will go for more complicated Humean accounts of the nature of chance. So long as the truthmakers for theories of chance are total histories, and so long as theories of chance are constituted by conditionals that assign chances to particular events on the basis of past history, theories of chance will not in general assign their Humean truthmakers a probability.²⁴

It is *agents* using theories of chance who have beliefs about which theories of chance are true, and who have to reconcile their beliefs about which theory of chance is true with probabilities they assign to futures that would—considered as total histories, for the Humean—threaten to conflict with those

beliefs. Conflict arises not at the level of chance functions, but at the level of the credence functions of agents who use them. Lewis was usually careful to keep it there.²⁵ Not for the reasons above, but because the link between the truthmakers for theories of chance and the events to which it assigns probabilities is made by the Humean Hypothesis, which is a philosophical thesis, quite extrinsic to theories of chance themselves. Lewis' argument showed that theories of chance assign positive probabilities to futures that would—according to the Humean Hypothesis—undermine their truth. And we have agents using theories of chance to determine credence. The difficulty concerns the interaction between the agent's doxastic commitment to the theories they employ and the probabilities assigned in the first-order content of the theories. And it puts pressure neither on the agents nor on the theories, but on the hypothesis that connects the truthmakers for those theories to the events to which they assign probabilities.

But we can get at the worry by considering the doxastic commitments of an agent like Cassandra, if we allow her beliefs both about total histories and whether she is the true expert on the panel, and suppose she is convinced with certainty of the Humean Hypothesis. Can we elicit a contradiction from her with the right series of questions? I don't see how. If we query her about her own expertise at any point in her career, she would say, "Nothing's ruled it out so far, but there's still some chance I will be shown wrong. We'll have to wait until all of the facts are in." She is infallible about the chance of fallibility conditional upon any finite body of evidence, and part of her wisdom consists in the recognition that, until the facts are in, there is always a non-zero chance of mistakes to come. Once we abandon the foolish idea that one has to conditionalize on one's doxastic infallibility every time one forms a first-order belief, we won't require her to assign a probability 1 to her own expertise. So long as she doesn't assign a probability of 1 to anything that entails her *non*-expertise, she will be consistent, and so long as she leaves open the possibility of more events to follow, she will never assign a probability of 1 to anything that entails her *non*-expertise.²⁶

Cassandra *would* get into trouble if she tried to apply the general recipe to herself, i.e., if she tried to use her less than perfect confidence in her own credence-assigning policy to downgrade the probabilities it generates. For then she would follow every assignment of probability with a second thought that assigned a new, lower probability. This is one of the puzzles generated by allowing beliefs about the reliability of one's epistemic policies. Circles of this kind can turn nasty for all kinds of beliefs, probabilistic or not, and they don't raise problems specifically for chances.²⁷

Suppose that we allow that Cassandra doesn't have to conditionalize on her own expertise in making a first-order assignment of probability. But what about a belief-forming agent using the advice of a modest expert to set his own credences? Such an agent does have to believe in the expertise of his chosen authority and he couldn't simultaneously have full confidence

in her expertise and accept her own assessment of it;²⁸ $Ch_{T_w}(T_w) < 1$ and hence $Ch_{T_w}(T_w) \neq Ch_{T_w}(T_w / T_w)$. That is correct. But no agent who is following the General Recipe will ever have full confidence in a theory of chance. He will always be drawing on evidence from a finite history, juggling multiple epistemic possibilities, never plunging full-stop for one. There is no contradiction here either. So, Lewis was right that if we assume the Humean Hypothesis, the credence function of an agent using **PP** in its original form would lead to the contradiction he diagnosed. But once we move from **PP**_{orig} to **PP**_{uncond}, and supplement it with the standard formula for reasoning under uncertainty, the contradiction is dispelled. No agent can consistently have full confidence in a theory of chance that recognizes the possibility of futures that undermine it, but no agent who is following the General Recipe will ever have full confidence in any theory of chance. He will be an inveterate fence-sitter, dividing his opinion between epistemically possible theories of chance, ruling out those only that come into conflict with the evidence.

Hall and Roberts

This proposal builds on two papers in the literature: a paper by Ned Hall, mentioned already above, and a recent, related paper by John Roberts. Hall's paper is a landmark in discussion of chance and its influence is deserved. Hall's paper develops the insight that consistency for **PP**_{orig} can be regained by a relatively simple manoeuvre. If you allow yourself to conditionalize on T_w , simply temper your credences by adjusting them to the chance conditional on T_w . The same ought to be said for *any* proposition; if $Ch_{T_w}(B) < 1$, your credence in a proposition A, given B shouldn't be the chance of A *simpliciter*, but it's chance conditional on B. I may believe that there's a very good chance A gets tenure given that she wins a Nobel prize, but that tells me nothing by itself about the degree of belief I should assign A's getting tenure. Only in the special case that it's certain, or close to certain, that B is true, can we translate beliefs about the chance of A, given B into credence in A. The revision of **PP**_{orig} that he proposed and that Lewis eventually accepted was:

$$\text{NP } Cr(A/H_{T_w}T_w) = Ch_{T_w}(A/T_w)$$

He argued that **PP**_{orig} is obtained from **NP** by approximation on the hypothesis that after any appreciable history, the correct theory of chance assigns itself a probability close to 1, i.e., that by now, at least, if T_w is the correct theory of chance, $Ch_{T_w}(T_w) = \text{almost } 1$. We allow agents unqualified beliefs about what the chances are, given the truth of one or another theory, and let them turn these into beliefs about the chances themselves only if they have something close to certainty about which theory of chance is true. This is already half-way to the General Recipe. He presumes we have independent

grounds for relative certainty about the correct theory of chance, and he seems clearly to have in mind a mixing procedure for determining credences where there is non-negligible uncertainty (Cassandra, as I mentioned, is his invention).

Roberts' paper, coming much later than Hall's, gives a kind of pragmatic vindication of the concept of chance. He accepts a form of **PP**_{orig} generalized to accommodate uncertainty about the correct theory of chance, and argues that the generalized principle leads to contradictions under some conditions (*viz.*, when one conditionalizes on one or another T_w), but since those are not conditions that epistemic agents like us (i.e., agents whose information about the world is always drawn from history) are expected to encounter, we will never get into trouble applying the principle. So Roberts' generalized principle —**GPP**, below —which is, he argues the one that actually governs the way we use information about chance in day to day reasoning, has the same status as a prescription of the form “follow the flow of traffic and never speed”, except that there are external contingencies that keep us from encountering the circumstances under which it can't be consistently followed (in this case, when the flow of traffic exceeds the speed limit).

$$\mathbf{GPP} \text{ Cr}(A/E) = \Sigma a\text{Ch}_{T_w}(A)$$

GPP is equivalent to the special case of the General Recipe, when E contains only historical information, i.e., when $H \rightarrow E$.

In one respect, Roberts' discussion is an advance on Hall. His **GPP** guides credence even under conditions of uncertainty about T_w . In another, however, it's a retrograde manoeuvre; it's inconsistent without restriction, and the restriction relies on a notion of admissibility that is given no clear, non-question-begging characterization. According to **GPP**, one is to let one's best estimate of the chance guide credence, except when one has ‘information about the future of the sort one gets from crystal balls’. Hall's NP is consistent and makes no use of an undefined notion of admissibility, but provides no guidance for unconditional credence where there is uncertainty about the correct theory of chance. Hall and Roberts both saw that **PP**_{orig} is consistent just in case $\text{Cr}(E) = 1$, where E is the conditional basis; Hall's response was to amend the principle as *per* NP by conditionalizing on T_w on the right. Roberts' was to generalize the principle and then restrict its application so that it is only applied under conditions of the sort that epistemic agents like us expect to encounter. **PP**_{uncond}, by contrast, is perfectly consistent, and can be followed always and everywhere. It would be a *pragmatic*, rather than a *logical* mistake to try to follow it if one is in possession of information from the future, and we can incorporate that into its conditions of application. The General Recipe derived from **PP**_{uncond} (by principles for accommodating uncertainty that have nothing specifically to do with chance) is an advance on Roberts because it derives Roberts' principle as a special case

under conditions in which it can be consistently applied, provides explicit characterization of inadmissible information (in the form of conditions under which it diverges from **GPP**), and derives the contradiction obtained by applying **GPP** in those cases directly from the definition of chance. It is an advance on Hall because it can be used to guide credence where there is uncertainty about the chances. Most importantly, it provides the basis for a physical interpretation for our beliefs about chance. It puts us in a position to identify features of the world *could*, and *should*, play the role it describes in guiding credence for epistemic agents of the sort we are.

Conclusion

To review; Lewis thought that the concept of chance was captured by a principle that counseled that if you know what the chances are, and you have no crystal balls or magical insight into the future, you should adopt the chance of an event *e* as your credence in *e*. The defeasibility of **PP** in the presence of information about the future was essential to its plausibility. The problem that Lewis saw for Humean accounts of chance was that such accounts entailed that the theory of chance itself contained information about the future, and since he thought any application of the principle had to include the theory of chance in the conditional base, that seemed to him to render the principle empty, rather like saying ‘adjust credence to chance except when the earth moves, the stars shine, or the sun comes up’. My resolution was simply to deny that the theory of chance should be included in the conditional base in **PP**. An agent with no magical sources of information about the future would understand that she should adjust her credence in *e* to the chance of *e* if she knows what the chance is and has no overriding information from crystal balls. Of course, if she has no overriding information from crystal balls quite generally, she won’t know with certainty what the chances *are*, because she won’t know which theory of chance is true. But so long as she knows what the various theories of chance say about the chances (i.e., so long as she can carry out the calculation that tells her what the chances are conditional on one or another theory of chance being true), and so long as she has epistemic probabilities defined over theories of chance, she can let her best current estimate of the chances determine her credence. Chance would be guiding credence, but in a way tempered by ignorance. That was the proposal summarized in the General Recipe.

Notes

¹I am extraordinarily grateful to two anonymous readers for *Nous* who went beyond the call of duty in making this a much better paper than it would have been if it was solely mine. The first I owe a beer and a handshake. The second owes me a good stiff drink. Many thanks to you both.

²Note here, that the Humean needn't have anything against primitive chancy facts, so long as those facts are local. What the Humean does object to is primitive connections between distinct facts. Facts about the chance of *e* in the here and now can't constrain how things are with *e*, and *vice versa*. This is a frequent misunderstanding.

³This can be translated into relativistic language by replacing 'past' and 'future' with 'events that fall, respectively, inside and outside the past light cone'. Events that are space-like separated from *p* have the same status, for these purposes, as events that fall inside the future light cone, and we include an event in its own past light cone.

⁴This is Lewis' stage-setting. Nothing hinges crucially on the details. Events are conceived here as concrete particulars individuated in part by spatiotemporal location. And to say that we can always ask 'what is the chance of *e* at *p*?' is not to say we will always get an answer. We can allow undefined values.

⁵There are good reasons not to follow Lewis here. All that is needed for what follows is that it acts as an implicit definition chance, identifying the target of investigation. Chance is that objective feature of the world that plays this role in guiding belief.

⁶This passage comes from "Humean Supervenience Debugged", *Mind*, 1994, p. 475. It recapitulates the definition in "A Subjectivist's Guide to Objective Chance".

⁷For details, see Albert, *Time and Chance*. Some of the details of Albert's account are disputed, but however it is filled in, it will provide a physical characterization of information from the past.

⁸The subscript denotes the world at which *T* holds; in all that follows, *w* will be the actual world.

⁹This is essential for the argument. Noone, to my knowledge, has tried to exploit the potential loophole here. Nor will I.

¹⁰"Humean Supervenience Debugged", *Mind*, 1994, p. 485. Thau's and Hall's papers are in the same volume. I have substituted my own notation for Lewis; 'Cr' stands for credence and 'Ch' for chance; both are presumed to be probability functions, which is to say simply that they provide an assignment of real numbers to events that satisfy the axioms of the probability calculus.

¹¹*Ibid.*, p. 487.

¹²More precisely, so that the facts about the world in virtue of which the *p*-chances are what they are, lie in *p*'s past light cone.

¹³The special case of the emended principle that makes it clear how it deals with undermining futures is NP: $C(A/HT) = \text{pr}(A/T)$. Hall's contribution to the same volume, more on which, below, derives it from the more general principle, and provides theoretical motivation.

¹⁴*Ibid.*

¹⁵By 'objective', I mean they can't be descriptive of belief in any obvious sense. This follows from their role in guiding belief.

¹⁶Even those following Lewis who explicitly make the distinction appear to suppose that we could move seamlessly from reasoning within the scope of a theory to reasoning *in propria persona*, where there is theoretical uncertainty, by simply conditionalizing on its truth (i.e., they suppose $\text{Ch}_{T_w}(A) = \text{Pr}_r(A/T_w)$), as when you move from parroting the opinions of your parents to speaking in your own voice by adding 'according to X' or 'supposing that what X says is correct'. The chance of *A* at *t* according to *Tw* is *not* in general the same as the probability of *A* conditional on *Tw*; *Tw* itself, if it is a modest theory, furnishes a counterexample -viz., $\text{Ch}_{T_w}(T_w/E) \neq \text{Pr}_r(T_w/ET_w)$, and there are countless others (substitute for *T_w* any proposition the probabilistic effect of *T_w* on which is not screened off by *E*).

¹⁷The difference between epistemic policies and beliefs is one that van Fraassen has recently argued is the key to arriving at a consistent, undogmatic empiricism (*The Empirical Stance*, (New Haven, Yale University Press, 2002)).

¹⁸Van Fraassen has recently used the distinction to respond to charges that belief in empiricism is self-undermining because it is not itself an empirically testable hypothesis by arguing that it should be regarded as an epistemic policy (or, in his words, 'stance') rather than a belief.

And he has given quite general reasons in *The Empirical Stance* and elsewhere for resisting the assimilation of the of the pragmatics of belief to their content.

¹⁹Notice that the general recipe is equivalent to a principle that we will see later that says adjust your *conditional* credence in (A*/E) to the weighted sum of probabilities assigned the conditional (A/E) by different theories of chance. **GPP_{ures}**: $Cr(A*/E) = \sum ChT_w (A*/E)$ We will see later a principle proposed by John Roberts that is a special case of **GPP_{ures}** that holds when E is purely historical information.

²⁰More precisely, the chances are those features of the world that bear a simple, singular relation to credence; *viz.*, they are those that guide the credences of reflective reasoners.

²¹The Principle, as stated, governs reasoning in the face of ignorance about the true theory of chance, but presupposes complete knowledge of history. A similar mixing procedure would need to be employed to accommodate ignorance of history, and ignorance of history will be expected to affect credence in theories of chance in the same way that ignorance of the future does. I think one of the reviewers for pointing out this lacuna.

²²The example comes from Hall, but on the notion of an ‘expert function’, and the idea of thinking of probabilities as credences of experts, see van Fraassen, *The Empirical Stance*, Yale University Press, 2004.

²³Adding the rest of history would leave the probability untouched; the present physical state of a system screens off the effect of anything that falls in its past light cone.

²⁴One can adopt definitions that support (or demand) the extension of chance functions to total histories, at the price of losing the immediate connection with physics, and at the price of obscuring the real source of the problem. Hall, for example, rejects Lewis’ identification of theories of chance with history-to-chance conditionals, and assumes that chance functions provide complete distributions over the space of total histories. It facilitates the technical exposition, but at the price of losing touch with physical examples, and, as argued below, obscuring the argument.

²⁵Note that it’s perfectly okay to talk of the probability that a theory of chance assigns to an undermining future, if we mean by this futures that would – considered as total histories – undermine the theories of chance which produce them. Lewis can be read this way.

²⁶This is exactly how a Humean should reason; there is an extension of any finite initial segment of history that would, considered as a total history, vindicate her expertise, as well as one that would undermine it, and anything more than a qualified doxastic commitment would be committing oneself to oracular knowledge.

²⁷Perhaps the most familiar instance of it is the preface paradox, in which an author expresses confidence that some of what follows is bound to be incorrect. You can form beliefs about the reliability of your belief-forming procedures, or the probability that all of your beliefs are correct, but they are inveterately meta-beliefs. They can’t be incorporated into the belief-forming policy itself. You can form beliefs about the reliability of your policies, but as soon as you use those beliefs to readjust your credences, you have abandoned the original policy.

²⁸For simplicity, I’m equating ‘believe’ or ‘believe without reservation’ with ‘assign probability 1 to’. This can be made more precise, but at a cost in exposition.