

VIII—EPISTEMIC DEFERENCE: THE CASE OF CHANCE

JAMES M. JOYCE

Epistemic deference is the phenomenon in which one person uses the deliverances of some information source, perhaps the opinions of another person, as a model for what to believe. The paper aims to clarify the nature of epistemic deference in probabilistic contexts, to explain the conditions under which deference is appropriate, and to examine deference to objective chances, as epitomized in David Lewis's Principal Principle. This latter analysis will show, in contrast with views that portray chance as an ideal inductive logician with total recall, that our deference to chance is grounded in contingent limitations on our ability to access information and our recognition that the physical probabilities that instantiate the actual chances codify all the types of information that humans are able to possess.

My broad topic is *epistemic deference*, the phenomenon in which one person uses the opinions of another, either a real person or some idealized information source, as a model for what to believe. Three questions will be addressed: How should we model epistemic deference within subjective probability theory?¹ What is it for a believer to *regard* another as being worthy of deference? Can we isolate substantive conditions under which one believer should defer to another? In addressing the last question I will be considering the case of *objective chance*, which is often used as a paradigm for deference to an *epistemic expert*. I will conclude that, contrary to views that portray chance as a kind of super inductive reasoner, the basis of our epistemic deference to chance is primarily to be found in facts about our limitations as knowers.

¹ I use the framework of subjective probability theory not because I think it psychologically realistic, but because it provides useful first step in evaluating claims about belief and justification. Claims in general epistemology that are implausible in the idealized probabilistic framework are unlikely to hold up in more realistic contexts.

I

A Framework for Talking About Evidence. Let us imagine items of evidence as information about the values of probability functions. In simple cases, these are propositions of the form $q(A)=x$, where q specifies a probability function non-rigidly, A is a proposition, and $0 \leq x \leq 1$ is a real number.² For example, if A says that the next Prime Minister of the United Kingdom will be from the Conservative Party, then I might have evidence of the following sorts:

$me_{now}(A) = 0.5 \leftrightarrow$ My current subjective probability for A is 0.5.

$ladbros_{now}(A) = 0.6 \leftrightarrow$ The bookmakers at Ladbrokes are now accepting bets for A at odds of 6–4 on.

$ch_{now}(A) = 0.4 \leftrightarrow$ The current objective chance of A is 0.4.

This is a more general way of representing evidence than one might imagine. Since *truth-value-at-world-w*—the function $truth_w(A)$ that assigns A a value of 1 (or 0) iff it is true (false) in w —is a probability function, this framework is adequate for any sort of information that can be expressed propositionally. In more general cases, a person's evidence might comprise data about values of probability functions over a range of propositions. Such data can be expressed as a (perhaps infinite) conjunction: $(q_1(A_1) = x_1 \wedge q_2(A_2) = x_2 \wedge q_3(A_3) = x_3 \wedge \dots)$. Or, it might be that a believer only knows that the various probabilities fall within various interval ranges. In other situations, a believer might have evidence about the relationship among the values of non-rigidly specified probability functions, as when one learns that the bookies at Ladbrokes and William Hill are

² Formally, we are dealing with a probability space defined in the following way: Start with a Boolean algebra of propositions Ω , and some finite set of probability functions $Q = \{q, r, s, \dots\}$ defined on Ω that are specified *non-rigidly*. First, we extend Ω to Ω' , the algebra generated by all propositions $(B \wedge q(A) = x)$ for $A, B \in \Omega, q \in Q$, and $x \in [0, 1]$. Proceeding by recursion, Ω^{m+1} is the algebra generated by all propositions $(B \wedge q(A) = x)$ for $A, B \in \Omega, q \in Q$, and $x \in [0, 1]$. Our probability space is $\Omega^* = \lim_m \Omega^m$, the smallest set containing all the Ω^m . It is crucial here to denote probabilities using non-rigid designators to avoid Miller's paradox (Miller, 1966). To keep things straight, I am using lower-case italicized letters when a probability is specified non-rigidly, and upper-case bold letters when it is specified rigidly. Thus, in ' $q(A) = P(A)$ ', $q(A)$ is a non-rigid description, like 'the bias of this coin landing head', and $P(A)$ is a real number, like 0.25.

quoting the same odds on some uncertain prospect. Here the data have the form $q_1(A_1) \leq q_2(A_2)$. In still other cases, a believer might know that some non-rigidly specified probability is identical to some rigidly specified one, as when the probability that gives the potential outcomes of the toss of a certain die is known to be the uniform distribution over the faces. Here we write the evidence as $q = \mathbf{P}$, on the understanding that this means $q(A) = \mathbf{P}(A)$ for all propositions A for which $q(A)$ makes sense.

I will assume a simple probabilistic model of learning on which acquiring the data that $q(A) = x$ involves updating one's opinions by Bayesian conditioning. A person whose 'prior' opinions are represented by a subjective probability, or *credence function*, C (always assumed to be coherent) will take the data $q(A) = x$ into account by adopting a new credence function $C(\bullet | q(A) = x)$ provided $C(q(A) = x) > 0$. When the evidence is that A is actually true we will write this as $C(\bullet | A)$.

II

Epistemic Dference. Believers can have attitudes of deference, indifference or disdain toward sources of probabilistic information. These attitudes are revealed in a believer's tendencies to shift her credences toward or away from the source's values when she discovers what these values are. Here are some of the attitudes a believer with credences C might have with regard to the information that x is q 's probability for A :³

No Dference: C might treat q 's probability for A as irrelevant to A 's truth or falsity, in which case $C(A | q(A) = x) = C(A)$.

Weak Dference: C might assign q some weight on the matter of A but not as much as she assigns herself, so that $C(A | q(A) = x)$ falls between $C(A)$ and $\frac{1}{2}C(A) + \frac{1}{2}x$.

Equal Dference: C might see q as an *epistemic equal in re A*, and so 'split the difference' by setting $C(A | q(A) = x) = \frac{1}{2}C(A) + \frac{1}{2}x$.

³ One can also have relations of super-deference, and anti-deference.

Strong Deference: C might see q 's values as a better guide to A 's truth-value than her own, but still assign her own views some weight, so that $C(A|q(A)=x)$ falls between x and $\frac{1}{2}C(A) + \frac{1}{2}x$.

Complete Deference: C might take q as what Gaifmann (1986) calls an *epistemic expert* when it comes to A , in which case $C(A|q(A)=x) = x$.

A person's deferential attitudes can depend on both the content of the proposition in question and the probability assigned to it. Even sources to which one defers without reservation on certain subjects, might merit disdain regarding others. Likewise, even for a fixed proposition, one might defer to a source's probability assignments when these fall in certain ranges, but reject them when they fall in others. For example, after a physical examination I might well defer to my physician's opinion on the probability of my living to see my hundredth birthday if he reports a value of less than 0.2. But, I might ignore him entirely if he announces a probability of 0.99. Here, I am in the situation Hume described in the 'On Miracles' section of the *Enquiry*, in which an otherwise reliable source might not be believed simply because its reports are too incredible. Quoting the ancient Roman saying, 'I should not believe such a story were it told me by Cato,' Hume concludes that the antecedent incredibility of an event can invalidate even a great authority's testimony to its occurrence. Whether or not we follow Hume in thinking that this undermines reports of miracles, he is surely right that the appropriateness of epistemic deference toward information sources often depends on what information those sources convey.

The probability that an information source assigns to A is only likely to influence one's deferential stance toward the source when one has strong independent reasons for assigning A a different probability. In most other circumstances—especially those in which one's opinion about A is based on flimsy evidence, or when one has a firm view about the source's reliability—one's esteem might not depend at all on its particular probability assignments. The strength of one's deference to the source can then be measured by a single parameter $\lambda \in [0, 1]$ that satisfies the equation $C(A|q(A)=x) = \lambda x + (1-\lambda) \cdot C(A)$ for every $x \in [0, 1]$. When someone defers to a source in this way we shall say that source serves as an *epistemic comrade in*

re A for her. Epistemic experts are comrades with $\lambda = 1$. Epistemic equals are comrades with $\lambda = 1/2$.

The deference implied in epistemic comradeship is always grounded in a believer's antecedent expectations about what a comrade's probabilities are likely to be. One can best see this by noting that the following are necessary conditions for C to regard q and r as epistemic comrades *in re* A :

Agreement in Expectation: $C(A) = \sum_x C(q(A) = x) \cdot x$.

Agreement on Certainties: $C(A) = 1$ iff $C(q(A) = 1) = 1$.

No Expected Divergence Among Comrades:

$$\sum_x C(q(A) = x \wedge r(A) = y) \cdot (x - y) = 0.^4$$

These requirements enforce a kind of agreement among epistemic comrades. While one can disagree with one's epistemic comrades about A , and can even regard this disagreement as quite likely, one must agree with one's comrades *in expectation*. That is, one's expected value for all of one's comrades' probabilities for A must coincide with one's credence in A .

III

What Makes It Reasonable to Defer? Nothing said so far tells us anything about when it is or is not reasonable for a believer to treat an information source as an epistemic comrade. Two related issues arise. First, it would be useful to have some inventory of the features that a source should possess to make it reasonable for a believer to treat it as a comrade. Second, though deference is largely a 'local' matter that depends on the particular character of a believer's evidence about a source, one wonders whether there are any especially worthy sources that are due deference by all, or at least most, rational believers.

In addressing the first issue, it will be useful to draw on an illumi-

⁴ Keep in mind that one can expect two probability functions to agree in this sense, and yet know that they will not in fact agree. Expectations are not beliefs.

nating discussion of expert probabilities by Ned Hall (2004). Hall suggests two reasons for us to regard q as an expert. First, q might be better informed than we are (though perhaps no better at analysing data), in which case Hall calls q a *database-expert*. Alternatively, q might be better than we are at evaluating the relevance of data (without possessing more of it), in which case q is an *analyst-expert*. Adam Elga (forthcoming) employs a similar distinction (using the term ‘guru’ for Hall’s analyst-expert):

When it comes to the weather, I completely defer to the opinions of my local weather forecaster ... conditional on her having probability x in any weather-proposition, my probability in that proposition is also x ... In treating my forecaster this way, I defer to her in two respects. First, I defer to her information: ‘As far as the weather goes,’ I think to myself, ‘she’s got all the information that I have—and more.’ Second, I defer to her judgment: I defer to the manner in which she forms opinions on the basis of her information.

Similarly, when Elga characterizes the notion of an *epistemic peer* (whom he believes should be treated as an equal, $\lambda = 1/2$) he says that we should regard someone as our epistemic peer when (a) they possess all the information we have, and (b) are as good as we are at evaluating the truth-values of claims in light of this shared evidence, so that, before evaluating a claim, we think that they and we are equally likely to adopt the correct belief about it in light of shared evidence. Similarly, Tom Kelly characterizes an epistemic peer this way:

Two individuals are **epistemic peers** with respect to some question if and only if:

- (i) they are equals with respect to their familiarity with the evidence and arguments which bear on that question, and
- (ii) they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness and freedom from bias. (Kelly, 2005, p. 75)

The connecting thread here is the idea that the degree of epistemic deference we owe an information source is a matter of both (a) the amount and character of the evidence that the source’s probability reflect, especially relative to our own evidence, and (b) the degree to which the source is better or worse than we are at arriving at accurate probability values on the basis of whatever evidence it is given.⁵

To sharpen (a) and (b), begin by noting that C’s tendency to defer

to q might stem from C 's belief that q 's probabilities are based on information she lacks. In the simplest cases, this situation can be characterized as follows:

C 's deference to q is based on a lack of information about E when: C is uncertain about E ($0 < C(E) < 1$); C is certain that q is certain about E ($C(q(E) \in \{0, 1\}) = 1$); becoming certain about E will make q 's probability for A irrelevant to C , so that $C(A|E) = C(A|E \wedge q(A) = x)$ and $C(A|\neg E) = C(A|\neg E \wedge q(A) = x)$ for $x \in [0, 1]$.⁶

Here, C 's deference to q does not require her to admire q 's prowess in reasoning: she defers to q solely because she believes q knows more than she does about E .

Even when we revere an information source for its reasoning powers, we still might not defer to it simply because we have information it lacks. Even if my friend Ian is a splendid political analyst, I might not defer to his judgement about the likely party of the next Prime Minister if I have just received insider information, which he lacks, about a brewing scandal that is likely to lead to the resignation of half the Conservative MPs in Parliament. On the other hand, if I really do respect Ian's political acumen, I will again defer to him as an expert once he has been apprised of the scandal. So, say that

C 's lack of deference to q depends on her view that q lacks the information E when: C is certain of E ; C believes q is unsure of E ($C(0 < q(E) < 1) = 1$); and C would defer to q if q knew E , so that the conditional probability $q(\bullet|E)$ is a comrade for C .⁷

C 's deference to q can also depend on q lacking certain evidence. Suppose Ian is a member of a messianic cult of Conservatives who

⁵ Deference does not, however, require a source to have the same evidence we have, nor to draw conclusions by the same inferential means we use. As a juror in a trial I might have lots of evidence that the eyewitness lacks, but still defer to her as to the suspect's guilt because her private evidence is so much better than mine. Likewise, if George's 'gut intuitions' about horse races tend to be very reliable, I might defer to him even though I would never trust my own 'gut intuitions' about such matters.

⁶ C 's deference to q in re A stems partly from a lack of information about E , when there is $\mu < \lambda$ such that $C(A|E \wedge q(A) = x) = \mu \cdot x + (1 - \mu) \cdot C(A|E)$ and $C(A|\neg E \wedge q(A) = x) = \mu \cdot x + (1 - \mu) \cdot C(A|\neg E)$ for $x \in [0, 1]$.

⁷ C 's lack of deference to q in re A stems partly from q 's lack of information about E , when there is a $\mu > \lambda$ such that $C(A|q(A|E) = x) = \mu \cdot x + (1 - \mu) \cdot C(A)$ for all $x \in [0, 1]$.

subscribe to a bizarre prophecy in which, sometime during the twenty-first century, a direct descendant of both Winston Churchill and Margaret Thatcher will lead the Conservatives to victory. Suppose further that, unbeknownst to Ian, we have just discovered that in 1957, when Thatcher was 32 and Churchill was 83, both stayed in the same hotel on the same night (which, let us assume, marginally raises the minuscule probability of the two having an illegitimate child, and so of there being a current direct descendant of both). Even if Ian had this extra information, however, we might not trust him to pick a winner in the next election. We might suspect that his messianic views, which do not typically influence his predictions, would be invigorated by information about a possible Thatcher-Churchill tryst, and that this might lead him to overestimate the Conservatives' chances.

This would not happen with a perfect information processor. Such a person will be worthy of deference no matter what evidence is placed before her. This is what Hall means by an 'analyst-expert'. If we imagine the totality of C's evidence given by some single proposition E_{tot} , which is believed certain, then

C defers to q as an analyst exactly if $q(\bullet|E_{tot})$ is a comrade for C relative to any body of potential evidence E_{tot} that C might possess.

On this picture, then, epistemic deference is a function of both what believers know about the data that information sources have at their disposal and the degree to which the probabilities issued by these sources reflect reasonable policies of drawing deductive and inductive inferences from data. This is not meant as a sharp distinction: there is no bright line between deference based on the quality of a source's information and deference based on the quality of its information processing. One's respect for a source as an analyst-expert will often depend on one's conviction that the source is well informed about key logical facts. Likewise, being a database-expert requires having a consistent set of data, and perhaps even having data that reflect facts about reliable information processing. Even if blunt, the database-analyst distinction is a useful tool for thinking about the broad kinds of issues that arise when questions of epistemic deference arise. In many concrete cases we can separate the

two factors, and can point to one or the other as the driving reason for deferring to a source.

This is especially apparent when we turn to our second question and inquire into the existence of sources to which deference is *universal* or very widespread. Consider $truth_{@}$, the function that assigns 1 to all actual truths and 0 to all actual falsehoods. The laws of probability ensure that every coherent believer must esteem $truth_{@}$ as an epistemic expert: $C(A|truth_{@}(A) = 1) = C(A|A) = 1$ whenever $C(truth_{@}(A) = 1) > 0$. $Truth_{@}$ does not need to be a great analyst to merit this respect. Because it has so much data, its reasoning strategy can be simple: it consults its premisses and asserts conclusions that are explicitly found therein, thereby begging the question. Deference to $truth_{@}$ is due to its exhaustive premiss set: it is a universal database-expert.⁸

Are there universal analyst-experts in the way that truth is a universal database-expert? Such a thing would be something like a perfect inductive reasoner who could come to the right conclusions no matter what evidence it is given. One way to be such an ideal reasoner, of course, is to know everything up front and ignore evidence. Truth has this feature, and so is a universal analyst-expert in virtue of being a universal database-expert. At the other extreme one might imagine an *a priori inductive* or *evidential probability*—*eprob*—which, despite lacking special access to empirical information, is always able to settle on the probability that is best supported in light of any given data. (I am not at all sure that such a probability exists, but will not air my doubts here.) If there is such a thing as an evidential probability function, then it will be a universal analyst-expert: all rational believers will defer to it whenever it is provided with their evidence.

Between $truth_{@}$, the universal database-expert, and *eprob*, the universal analyst-expert, there might be intermediate cases, in which we defer to some information source partly because it knows things we do not and partly because it is better than we are at drawing conclusions from the data. *Objective chance* may be something like this.

⁸ It is a universal analyst-expert as well. However, this is only because, for any coherent C and E with $C(E) > 0$, $truth_{@}$ is database-expert for $C(\bullet|E)$. So, while the database-analyst distinction is not sharp here, it seems clear that deference is underwritten by database considerations.

IV

Chance and the Principal Principle. Much of the philosophical discussion of objective chance stems from David Lewis's justly celebrated 'A Subjectivist Guide to Objective Chance' (1980). This article has, unfortunately, fostered a picture of chance's epistemic status that is misleading in two respects. First, Lewis has made it seem that the important philosophical issues about objective chance largely concern the fact that chance can undermine its own status as an epistemic expert. Second, by tying his theory about the epistemic role of chance to a specific sort of 'best system' metaphysics for chance, Lewis's work has tended to suggest that objective chance functions like a perfect information processor, and that this is what explains our deference to it. This picture, I will argue, obscures crucial facts about chance's epistemology. In particular, it obscures the fact that the probability functions that realize objective chances encode so much information about the causes of events that our epistemic limitations force us to see them as database-experts.

The epistemic role of chance, according to Lewis (1980, p. 87), is given by the *Principal Principle*, which says that the function that gives the chances of events at a time should be an epistemic expert for any coherent subjective probability function C that contains no information that is 'inadmissible' at that time.

- (PP) Let C be the rational credence function for someone whose evidence is limited to information about the history of the world up to time t . If E is admissible at t and if $C(E \wedge \text{chance}_t(A) = x) > 0$, then $C(A | E \wedge \text{chance}_t(A) = x) = x$.

Lewis saw this as a kind of definition of the concept of objective chance. 'A feature of Reality deserves the name of chance,' he wrote, 'to the extent that it occupies the definitive role of chance; and occupying the role means obeying [PP], applied as if information about present chances, and the complete theory of chance, were perfectly admissible' (Lewis, 1994, p. 489). Chance is thus defined by its special epistemic role: chance (i) functions as an epistemic expert for all believers whose evidence is restricted to information about the past and present, and (ii) fails to function as an epistemic expert *in re* A for believers who somehow possess information

about the future that is relevant to A 's truth or falsity. Anything that fits the bill (or gets close enough) is chance.

The limitation on C and the admissibility constraint are meant to ensure that C does not incorporate information about A 's truth-value that is not reflected in its chance at t .⁹ As Lewis (1980, p. 92) explains it, '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.' He also tells us that information about the history of the world up to t is always admissible, and that A and information about the future effects of A 's truth are generally inadmissible. And in Lewis (1980) he suggests that facts about current and past chances are always currently admissible.

As is well known, on Lewis's preferred 'Humean' theory of chance, facts about chance turn out to be *inadmissible* because they covertly convey information about the future. Lewis believed that the chances at any time t are determined by the world's history (past and future) of the particular non-chance events. Moreover, the identity of the chance function, in Lewis's 'best system' picture, is determined by considerations of fit with the observed frequencies, simplicity, and strength. (Notice that these are the same sorts of broadly inductive considerations that govern theory choice in general, and which we use when making inferences about the future based on information about the past.) The problem for Lewis was that his theory allows 'undermining futures' that have a non-zero chance of occurring, so that $\text{chance}_t(F) = x > 0$, but whose occurrence would preclude the chances from being what they are, so that $(F \wedge \text{chance}_t(F) = x)$ is certainly false. PP then requires that $C(F | \text{chance}_t(F) = x) = x > 0$, but logic requires that $C(F | \text{chance}_t(F) = x) = 0$.

With the help of Hall (1994) and Michael Thau (1994), Lewis (1994, p. 487) ended up solving the problem as follows. The core of the difficulty is that Lewis's theory of chance allows for the possibility that the current chances currently have some chance of being

⁹ In general, E is *admissible* (relative to q , A and C) just when adding it to C 's evidence does nothing to undercut q 's status as an expert about A for C . In particular, if q is an expert for C , then q remains an expert for $C(\bullet | E)$. Likewise, if E is *inadmissible*, then learning E does undercut q 's status. Specifying which propositions are admissible and inadmissible is critical to grasping a source's epistemic status. In general, the better a source, the more things that are admissible.

other than they actually are, so that both $chance_t(A)=x$ and $chance_t(chance_t(A)=x)<1$. If this is so, then it seems wrong to set one's credence (at t) for A at x when one supposes that A 's chance (at t) is x . After all, one is then supposing that one has information that the chance function itself lacks, namely the information that $chance_t(A)=x$. The natural solution is to modify PP so that it does not require you to defer to chance as an expert *tout court*, but as an analyst-expert, by deferring to it conditional on the information you possess. This yields the 'New Principle' (NP):

(NP) Let C be the credence function for someone whose evidence is limited to the past and present. Then, if the chances are given by probability function P ,¹⁰ then $C(A|chance=P)=P(A|chance=P)$. As a special case we have $C(A|chance_t(A)=x \wedge chance_t(A|chance_t(A)=x)=y)=y$.

NP has a lot going for it. The underlying epistemological idea is sound: as we have seen, when we have evidence that some generally reliable source lacks, we often defer to that source once it has been provided with our evidence. It is no different here: the information that the chance function lacks is just that its values are objective chances. As soon as chance is made aware of this it will provide answers to which we may defer. Second, NP does not rely on an admissibility restriction. When we acquire information that would have been inadmissible by the lights of PP, NP merely has us provide that information to chance and defer to its better judgement.¹¹ This seems like an advantage.

Hall (2004, p. 101), impressed with this aspect of NP, proposes that '*chance is an analyst-expert ... this claim holds for chances at any time, and without qualification*' (i.e. for any proposition A for which $chance(A|E)$ is defined). Hall's proposal, then, is this:

Chance is a universal analyst-expert: If C is a rational credence function, then for any body of evidence E with $C(E \wedge chance_t$

¹⁰ Here both *chance* and P should be thought of a two-place functions, so that this identity means that $chance_t(A)=P_t(A)$ for all A and t .

¹¹ Suppose we ask what we should believe if we somehow came to know A 's chance at some future time $t^*>t$. NP gives the right answer as long as (as is reasonable to assume) future chances trump current chances, so that $C(chance_{t^*}(A|chance_{t^*}(A)=x)=x)=1$.

$(A|E)=x) > 0$, one has $C(A|E \wedge \text{chance}_t(A|E)=x) = x$.

This is chance as ideal inductive reasoner: give chance any evidential input E , and it will spit out a probability to which a person with E as her total evidence should defer.

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

Chance is a database-expert about the past and present: If E is entirely about the world's state up to time t , then $C(E \wedge \text{chance}_t(E) < 1) = 0$, and $C(A|E \wedge \text{chance}_t(A)=x) = x$.

These two principles make for a natural interpretation of chance as it has come to be understood by philosophers in light of Lewis's writings. This is especially so if one thinks of chances as being fixed by best-system considerations (though one need not be committed to a 'best system' account of laws to find the principles appealing). They also offer an explanation for our deference to chance: as long as our evidence is no better than the information chance has at its disposal, and in so far as we recognize that we are not as expert as chance at drawing conclusions, we should defer to chance.

Still, one has the sense that something crucial has been left out. All the focus on correcting for undermining futures has directed attention away from a central feature of chance's epistemic role. Pic-

turing chance as a perfectly effective reasoner with perfect information about the past obscures the fact that our deference to the particular probability functions that actually realize chance is ultimately based on our views about the causal structure of the world, and about the sorts of restrictions that this structure places on the evidence we can possess.

One cannot see this clearly when one focuses exclusively on PP and NP, and takes these to be definitions or quasi-definitions of the concept of objective chance. Instead, one must look to the sorts of probability functions that realize the concept in the world we inhabit. Deference principles like PP and NP (non-rigidly) describe probabilities that can be realized in different ways in different worlds. It is not merely that there are a variety of specific probability functions P (rigidly specified) for which $chance_t = P$ might hold. Rather, chance might be realized by some particular type of probability function so that $chance_t = q$, where q is non-rigidly described, and it might be that our reason for deferring to chance in the way PP or NP require is that we believe this identity and defer to q in the way PP or NP require. Indeed, this seems to be the case. Chances in our world are realized by physical probabilities, like those found in quantum mechanics. Our reasons for deferring to these physical probabilities turn out to have much more to do with ‘database’ considerations than with ‘analyst’ considerations. Moreover, the possibility of undermining futures plays no role at all in the process.

Consider a concrete case. If you convince me, say, that you have correctly described the complete quantum-mechanical wave function for an isolated polonium atom, and if you also convince me that the probability of an atom with that wave function decaying within 138.876 days is $\frac{1}{2}$, then my credence for that event will be $\frac{1}{2}$. My reasoning is simple:

- (A) The quantum state of the system at any time encodes all physical information about the system at that time.
- (B) The quantum state of the system at any time acts a kind of ‘causal screen’ that separates past states from future events: no future event can be an effect of a past state except by being an effect of the present state—there are no causal pathways to the future that do not pass through the

present—and the present node in the causal path encompasses everything that is now causally relevant to the effect.

- (C) The quantum mechanical probability for decay reflects the totality of the information contained in the current quantum state that is causally relevant to the decay. It is a kind of summary statistic that measures the ‘causal tendency’ of the current state to produce the effect under consideration.
- (D) While there may be information that could undermine these probabilities (specifically information about the future), there is also a physical restriction that prevents us from ever having access to such information as evidence—we can condition on it (in the abstract), but cannot ever learn it. The restriction is that we, being physical systems, lack direct access to any information that is not about the present and the past—our evidence is confined to facts about our past light cone.

If one accepts these premisses (and one might question some of them: see below), then three factors make it reasonable to defer to physical probabilities when forming beliefs about the future.¹² First, because the current physical state of the system acts as a causal screen for past states, any causal influence that a past state can have on future events must be reflected in the current state. Second, since the present physical probability of an event encodes all information causally relevant to its occurrence that can be found in the present state, to know this probability is to know all there is to know about the present state as it pertains to the causes of the event.¹³ Third, since, as a matter of contingent fact, we cannot have direct evidence of the future, knowledge of the present physical state is the best evidence we can possibly have about any future event. Putting this all together, it should be clear why we should defer to physical proba-

¹² I believe similar reasoning applies to the probabilities in statistical mechanics (even though the underlying system is deterministic), and to standard coin-toss and dartboard cases in virtue of considerations like those described in Keller (1986). In both cases the important point is that there are physical restrictions on our ability to access information that would undermine the authority of the relevant physical probability function. This suggests that our reasons for deferring to quantum-mechanical probabilities are not centrally located in the fact that they describe indeterministic events.

bilities like those found in quantum mechanics. When probabilities reflect the causal powers of a state that is detailed enough to act as a causal screen, they serve primarily as *database-experts* for us: we defer to them partly because we are not able to acquire any relevant data about the causes of events that they do not already have. Of course, we also believe that quantum mechanics spits out the correct probabilities given the information encoded in quantum states, and so there is an ‘analyst’ element in our deference as well. However, we do not think of these probabilities as universal analyst-experts in the way NP requires. There are many items of potential information (e.g. the statement that quantum mechanics is false) about which quantum mechanics cannot speak.

Some observations and caveats. Notice first that physical probabilities need not be universal database-experts about the past. Information may be lost in the transition from past states to the present state, but none of it is causally relevant to the future except in so far as it left marks on the present. Second, nothing in this picture requires the current chances to supervene on the present, past or future pattern of non-chancy facts. Third, it is essential to this picture that the expert status of physical probability derives partly from general limits, albeit contingent ones, on our ability to acquire evidence of certain types. If we could directly intuit the future, or if we lived in a world where backward causation permitted reliable effect-to-cause inferences, we would have no use for physical probabilities, and their values would not be objects of deference for us. The expert status of physical probability derives, in large measure, from our own epistemic limitations.

This last point highlights a key difference between the account of *chance* that one finds in Hall or Lewis and the one mooted here. On the current proposal, there is no need to think of *chance* on analogy with a perfect inductive reasoner, or to portray it as an analyst-expert of some general type. Such ideas are motivated by concerns with the failure of PP in the face of undermining futures, and these

¹³ It can be instructive to think about physical probabilities as they evolve over time. It is incorrect to think of probability values at one time q_t as *causing* values at a later time q_{t+1} . But because these probabilities do track the underlying causal facts, the epistemology proceeds exactly as it would if the relation were causal. Knowledge of q_{t+1} screens off facts about q_t in much the same way that knowledge of an intervening cause screens off knowledge of any past cause that helped produce it (and is not linked to the effect by an independent causal pathway).

concerns are irrelevant to our reasons for deferring to the physical probabilities that actually realize *chance*. Failures of PP would be problematic if they involved inadmissible information *we could learn* and that could serve as evidence for drawing conclusions about the future. As quantum probabilities illustrate, this need not be so. Cases where *chance's* expert status is undercut by information we cannot possess are, on the present proposal, treated as 'don't cares' because they are 'can't knows'. Think, for example, of undermining futures. Suppose that the current physical probabilities are such that $q(F|q(F)=x)=y$. This is a bizarre situation; it means that the causal facts are such that if the causal facts are such that F will occur with probability x , then the causal facts are such that F will occur with some different probability y . It is not clear how this could even be possible; what is clear is that, even if it is possible, we could never know that the causal facts are such that F will occur with probability x . $q(F)=x$ is the kind of proposition that can be conditioned on, but not the kind that can be learned, if the physical probabilities are as described. So, undermining futures fall in the 'don't care' category even if they exist.

So, are these physical probabilities to which we defer in this way properly called 'chances'? To some extent, this is a question of semantics (in the pejorative sense). One might decide to reserve the honorific 'chance' for some super-probability that does know everything about the past and that never makes mistakes in reasoning. If so, then the sorts of physical probabilities I have been describing are not chances. But what's in a name? There are three things that I would claim for the sorts of physical probabilities to which people tend to defer: (i) they function just like chances for people with the sorts of limited evidence that human beings have; (ii) our patterns of deference to them are explained by their ability to reflect causal relationships; and (iii) our sense that *chance* is worthy of epistemic deference derives from these considerations, and not from any sense that *chance* is somehow especially expert at processing information, including information to which we have no epistemic access.

Let me close by considering a few objections that one might be inclined to raise against the little argument in (A)–(D).

Objection: What if (A) is false? What if there are 'hidden variables' on which the quantum states and probabilities depend? If

there are, then quantum states do not act as causal screens for all past states. In particular, they do not act as causal screens for past hidden variable states. But, if H is a past hidden state and A is an event for which the current quantum mechanical probability is $q(A) = x$, then q will not be an expert for anyone who knows H since, assuming that hidden variables obey deterministic laws, $C(A|H \wedge q(A) = x) = \{0, 1\}$.

Response: This does not matter so long as we cannot learn things about the hidden variables. If there are hidden variables, then the current QM probability for an event does not render all evidence about the past irrelevant, it just renders *all* the evidence about the past that we can actually have irrelevant. Not all information about the past is admissible, but all the information we can have about the past is admissible. Perhaps if one wants to insist that there are no chances in deterministic worlds, one may not want to call quantum mechanical probabilities ‘chances’. Even so, the reasons for deferring to quantum-mechanical probabilities would be the same.

Objection: What if (B) is false? What if there are ‘wormholes’ in space-time that allow objects and energy to traverse causal pathways from the past to the future without making any impression on the present? If this happens, the present state of a system would no longer act as a causal screen: some future events will be effects of past events without being effects of any present events.

Response: Here we *rightly* decrease our level of deference to quantum-mechanical probabilities. There are two ways this might happen. First, it could be that certain past events (the mouse going down the wormhole yesterday) are reliably tied to certain future events that lack present causes (the mouse suddenly appearing on top of the cheese tomorrow). If we had information about the past ends of these sorts of sequences, we would not defer to quantum-mechanical probabilities. We would, instead, seek a new sort of physical probability that agrees with the quantum probabilities when no ‘mouse going down the wormhole’ information is available, but which functions as an analyst-expert when such information is present. On the other hand, if there were no predictable ties

between events that ‘skip over the present’, then the quantum-mechanical probabilities might well remain worth significant deference, albeit without having expert status—a comrade of high rank.

Objection: Doesn’t this picture just assume that knowledge of causes screens off statistical information about the past?

Response: Yes, physical probabilities serve as evidential screens because they are associated with states that serve as causal screens. This is entirely proper. Physical probabilities, and anything else that comes close to chances in the world we inhabit, deserve deference because they reflect information about the potential causes of future events, and such causal information does screen off information about the past.¹⁴

*Department of Philosophy
University of Michigan
Ann Arbor, MI 48109-1003
jjoyce@umich.edu*

REFERENCES

- Elga, A. forthcoming: ‘Reflection and Disagreement’, *Noûs*.
- Gaifman, H. 1986: ‘A Theory of Higher Order Probabilities’, in J. Y. Halpern (ed.), *Theoretical Aspects of Reasoning About Knowledge: Proceedings of the 1986 Conference, Monterey, California* (Los Altos, CA: Morgan Kaufmann), pp. 275–92.
- Hall, N. 1994: ‘Correcting the Guide to Objective Chance’, *Mind*, 103, pp. 505–17.
- 2004: ‘Two Mistakes About Credence and Chance’, *Australasian Journal of Philosophy*, 82, pp. 93–111.
- Keller, J. B. 1986: ‘The Probability of Heads’, *American Mathematical Monthly*, 93(3), pp. 191–7.
- Kelly, T. 2005: ‘The Epistemic Significance of Disagreement’, *Oxford Studies in Epistemology*, 1, pp. 167–96.
- Lewis, D. 1980: ‘A Subjectivist’s Guide to Objective Chance’, in R. C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, Vol. II (Berkeley

¹⁴ For helpful discussion I thank Richard Bradley, Mark Kalderon, Boris Kment, Ian Rumfitt and Scott Sturgeon.

- and Los Angeles, CA: University of California Press); reprinted in D. Lewis, *Philosophical Papers*, Vol. II (Oxford: Oxford University Press, 1986), pp. 83–113. Page references in the text are to the reprint.
- 1994: 'Humean Supervenience Debugged', *Mind*, 103, pp. 473–90.
- Miller, D. 1966: 'A Paradox of Information', *British Journal for the Philosophy of Science*, 17, pp. 59–61.
- Thau, M. 1994: 'Undermining and Admissibility', *Mind*, 103, pp. 491–503.