

An Anti-Reductionist’s Guide to Evidential Support

Agustín Rayo

January 22, 2014

This paper develops an account of evidential support. The account is *anti-reductionist*: it does not presuppose that the notion of evidential support is reducible to ‘more fundamental’ notions. It is also *localist*: it does not presuppose that the notion of evidential support is well-defined independently of substantial background assumptions.

1 Reductionism and Ant-Reductionism

What sorts of properties should we use in building our conception of the way the world is? David Lewis offers one answer to this question; Robert Stalnaker offers another.

Lewis suggests a reductionist picture, according to which we should strive to build our conception of the way the world is on the basis of *fundamental* (or ‘perfectly natural’) properties.¹ Fundamental properties, according to Lewis, “render their instances perfectly similar in some respect. They are intrinsic; and all other intrinsic properties supervene on them.” Lewis is committed, moreover, to a supervenience thesis. He thinks that fundamental properties “figure in a minimal basis on which all else supervenes. No two possible worlds just alike in their patterns of instantiation of fundamental properties could differ in any other way.”²

¹See especially (Lewis 1983), (Lewis 1986*b*, introduction) and (Lewis 2009). Citations in this paragraph are taken from the latter.

²In the introduction to (Lewis 1986*b*), Lewis states a stronger version of the supervenience principle, and describes it as contingent. According to the stronger principle, no two possible worlds just alike in their *spatio-temporal* distributions of *point-sized* instantiations of fundamental properties could differ in any other way. The weaker principle I cite above appears in (Lewis 2009), and there is no indication that Lewis took it to be contingent.

Lewis's conception of fundamental properties is closely linked to his conception of modality. Because he thinks that there are no necessary connections between different instantiations of fundamental properties, he accepts a 'combinatorial principle', according to which one can vary the pattern of instantiation of fundamental properties across a possible world in arbitrary ways, and still get a possible world (Lewis 1986*a*, 2009).

Another feature of Lewis's conception of modality is that he thinks that counterfactual truths are ultimately made true by the intrinsic character of the actual world. This is because he thinks that what it takes for a counterfactual ' $A \Box \rightarrow C$ ' to be true is for all of the 'closest' A -worlds to be C -worlds, and he thinks that "it is the [intrinsic] character of our world that makes some A -worlds closer to it than others. So, after all, it's the [intrinsic] character of our world that makes the counterfactual true." (Lewis 1986*a*, p. 22) And, of course, Lewis's supervenience thesis entails that the intrinsic character of a world is fully determined by the pattern of instantiation of fundamental properties across that world. So, at the end of the day, counterfactual truths about our world are made true by the particular pattern of instantiation of fundamental properties across our world.

But then why bother talking about counterfactuals at all? The answer, Lewis thinks, is that "it is only by bringing other worlds into the story that we can say in any concise way what [intrinsic] character it takes to make the counterfactual true. The other worlds provide a *frame of reference* whereby we can characterize our world." (Lewis 1986*a*, p. 22, my emphasis)

Stalnaker contrasts Lewis's reductionist picture with an anti-reductionist one. He agrees with Lewis that other worlds provide a frame of reference with which to characterize our world. But unlike Lewis, Stalnaker thinks that "the possible-worlds frame of reference is playing an *indispensable role* in forming the concepts by which we describe our world at the most fundamental level." (Stalnaker typescript, my emphasis)

Let us consider the property of negative charge as an example. In order for it to count as a fundamental property in Lewis's sense, it would have to be *fully intrinsic*: whether

or not it was instantiated would have to depend entirely on the ‘qualitative character’ of its possible instances. And it would have to be capable of playing a different theoretical role than the one it actually plays. Suppose, for example, that F -ness is a fundamental property distinct from negative charge, which is instantiated in the actual world. Lewis’s combinatorial principle entails that there is a possible world w which is just like the actual world except that every instantiation of one of the two properties is replaced by an instantiation of the other. So if the best theory of F -ness for the actual world tells us that F s are disposed to attract one another, then the best theory of negative charge for w will tell us that particles with negative charge are disposed to attract one another (rather than repel one another, as they do in the actual world).

From a Stalnakerian perspective, in contrast, one can take the view that *part of what it is* to have negative charge is to be disposed to repel other particles with negative charge. An immediate consequence of this view is that there can be no world such as w , and, more generally, that the theoretical role of negative charge is fixed by its modal profile. Another consequence of the view is that negative charge is *not* a fully intrinsic property: whether or not one has negative charge is partly to do with whether or not one has the (relational) property of being disposed to repel other particles with negative charge.

Suppose this is right. In virtue of what do negatively charged particles have the relevant relational properties, and in virtue of what do they have the relevant modal profile? From an anti-reductionist perspective, the answer is that they *just do*: having those relational properties and having that modal profile is *part of what it is* to be negatively charged. There are no further, more fundamental properties that might be used to give a reduction of the relational to the intrinsic, or a reduction of modal facts to the intrinsic character of possible worlds.

A consequence of this anti-reductionist picture is that “one cannot separate the task of describing the world from the task of characterizing the space of worlds and the way

our world is related to others.” (Stalnaker typescript) When we describe the world as containing negatively charged particles, for example, part of what we’re saying is that it is *impossible* for those particles not to repel each other. So our description of the world is, in part, a characterization of the space of possible worlds. And there is no separating the two, since there are no further, more fundamental properties that might be used to decompose the property of negative charge. There is no set of purely intrinsic properties on the basis of which all else supervenes.

Lewis’s reductionism suffers from a shortcoming that Stalnaker’s anti-reductionism manages to escape. We have seen that Lewis’s combinatorial principle entails that fundamental properties are not essentially linked to their actual theoretical roles. As Stalnaker puts it, “the metaphysical theory underlying Lewis’s Humean project seems to require a radical gulf between the essential nature of a fundamental property—its quiddity—and the ways that the property manifests itself in the world.” (Stalnaker typescript) But since our cognitive access to a fundamental property is inevitably mediated by the ways in which it manifests itself in the world, this means that the essential nature of Lewis’s fundamental properties must forever remain hidden from us. We are, in effect, building our conception of the world on the basis of properties that are inevitably beyond our cognitive reach.³

On Stalnaker’s anti-reductionist picture, in contrast, there is nothing to stop us from building our conception of the way the world is on the basis of properties that are essentially linked to their theoretical roles. So the task of identifying the relevant properties is no different from the task of developing our best theories of the way the world is. In Stalnaker’s terms, “The development of a theory for predicting and explaining the phenomena involves the formation of the concepts by which the phenomena are described, and the testing of a theory against the evidence is also the testing of the legitimacy of the concepts that are formed.” (Stalnaker typescript)

³Lewis is explicit about this, and terms the predicament ‘Ramseyian Humility’. See (Lewis 2009).

1.1 Lewis on Objective Chance

The main aim of this paper is to develop an ant-reductionist account of evidential support. It will be useful, however, to start by considering the prospects of an anti-reductionist account of objective chance.

Because of the supervenience principle we discussed earlier, Lewis is committed to a *reductionist* account of objective chance: he is committed to the idea that the objective chances of a world supervene on the pattern of instantiation of fundamental properties across that world. So—unless objective-chance properties are to count as fundamental—he needs a *reductive* account of objective chance. He needs a story about what the objective chances *consist in* which is compatible with the idea that they ultimately supervene on patterns of instantiation of fundamental properties.

Frequentism is one such story. According to the most straightforward version of frequentism, the objective chance that an event of type E will yield outcome O is simply the *percentage* of (past, present and future) events of type E that yield outcome O . Thus, if 49.99983% of all past, present and future coin tosses land Heads, the objective chance that a particular coin toss will land Heads is exactly 0.4999983.

So understood, frequentism entails that there cannot be a world in which fair coins always land Heads. For if coins of a certain type were to always land Heads, then, according to frequentism, the chance of their landing Heads would have always been 1. So would have been the case that they were never fair.

Although Lewis does not endorse frequentism, he argues for an account of objective chance that yields a similar result. According to his best-system analysis of objective chance, “the chances are what the probabilistic laws of the best system say they are” Lewis (1994). In order to understand the proposal, we need to know what a *system* is, what a *best* system is, and what would count as a *law* of a system. A *system* for w is a set of axioms that truly describe w . The *best* system for w is the system for w that

delivers the best balance of simplicity, strength and *fit*, where a system is said to deliver good ‘fit’ to the extent that it avoids assigning low chances to events that actually take place. The *laws* of a system are the theorems of the system that describe regularities.

To illustrate the proposal, consider Seaborgium 265, which has a half-life of 8.9 seconds (i.e. the objective chance that a ^{265}Sg particle will decay within the next 8.9 seconds is $1/2$.) Let a ‘coin’ be a particular particle of ^{265}Sg , and let a ‘toss’ of that coin be the result of waiting 8.9 seconds; if the particle has decayed by the time the wait is over, we shall say that our coin toss landed ‘Tails’; otherwise we shall say that it landed ‘Heads’. Thus understood, coin tosses are genuinely chancy events. And they are always ‘fair’: every coin toss has a $1/2$ chance of landing Tails and a $1/2$ chance of landing Heads.

As in the case of frequentism, Lewis’s best-system analysis entails that there cannot be a world in which fair coins happen to always land Heads. For if a world’s ^{265}Sg particles never decay (i.e. if its coins always land Heads), then one would expect the best—simplest, strongest, best-fitting—system for that world to say that that ^{265}Sg is not radioactive, and therefore that the chance of Heads is 1. So the coin tosses of w were never fair.

Lewis was clear about about this point, and was inclined to put it by saying that although there is a sense in which it could happen that fair coins always land Heads, there is also a sense in which it couldn’t: “It could, in the sense that there’s a non-zero present chance of it. It couldn’t, in the sense that its coming to pass contradicts the truth about present chances. If it came to pass, the truth about present chances would be different.” (Lewis 1994, discussing a similar case).

Lewis was prepared to concede that this result was “very peculiar”, but eventually came to the view that it was not peculiar enough to give up the proposal (Lewis 1994). Indeed, he went on to note that his reductive analysis of objective chance renders the Principal Principle useless—a principle which he had earlier thought to “capture all we know about chance” (Lewis 1980). But instead of taking the result to cast doubt on the

reductionist proposal, he saw it as a reason to revise the Principle Principal.

It seems to me that this was a mistake. I think the Principal Principle needs no revising, and that it is the reductionist account of objective chance that must go. I think the reductionist account distorts the notion of objective chance by taking objective chances to have less counterfactual robustness than they ought to have. Contrast the property of being fragile with the property of being a future lottery winner. The property of being fragile is *robust*: if you're fragile you might break, but if you did brake it wouldn't have thereby been the case that you were never fragile. The property of being a future lottery winner, on the other hand, is *not* robust: if you are, in fact, a future lottery winner, it is possible for your to fail to buy a ticket, but if you did fail to buy a ticket it would have thereby been the case that you were never a future lottery winner. It seems to me that the property of having a certain objective chance is robust: although Lewis is right to think that it is possible for the world's fair coins to happen to always land Heads, he is wrong to think that if they had indeed always landed Heads they would have never been fair.

1.2 An Anti-Reductionist Account of Objective Chance

I would now like to sketch an anti-reductionist alternative to Lewis's account of objective chance. On the proposal I will be defending, there is no presupposition that the notion of objective chance can be spelled out in more fundamental terms. Our quest for fundamentality must stop somewhere; I see no reason think that it couldn't stop at the objective chances.

The best way of getting a handle on the proposal is to consider an example. Suppose that after a rigorous scientific investigation of the quantum world, we reach the conclusion that ^{265}Sg isotopes have a half-life of 8.9 seconds; we conclude, in other words, that the objective chance that a ^{265}Sg particle will decay within the next 8.9 seconds is $1/2$. In reaching this conclusion, we simply deployed our best scientific methods. Maybe our

best scientific methods are those that secure the best balance of simplicity, strength and fit; maybe not. Either way, there is nothing philosophically contentious going on: our conclusion about the half-life of ^{265}Sg is simply an instance of ordinary scientific practice.

Now suppose that we go on to claim, on the basis of our scientific conclusions, that *part of what it is* to be a particle of ^{265}Sg it to have a half-life of 8.9 seconds. In making this further claim, we proceed in the same sort of way as we would have if we claimed, on the basis of a scientific investigation of the quantum world, that *part of what it is* to be negatively charged is to be disposed to repel other negatively charged particles. In both cases we use the role that a property plays in our best scientific theorizing to identify the property's essential properties.

I do not wish to suggest that it is in any way *obvious* that these are the right lessons to draw from our best scientific theorizing. I do claim, however, that a Stalnakerian anti-reductionist framework makes room for such lessons to be drawn. Both in the case of ^{265}Sg and in the case of negative charge, we use our best understanding of the underlying science to characterize a property with which to build our conception of the way the world is. Neither of the resulting properties—being a particle of ^{265}Sg , and having negative charge—would count as a fundamental property from a Lewisian point of view, since neither of them is purely intrinsic. What we have instead are properties essentially linked to their theoretical roles, and an anti-reductionist might think that our conception of the way world is should be built from just such properties.

Let us now suppose that having concluded that part of what it is to be a particle of ^{265}Sg is for it to have a half-life of 8.9 seconds, we go on to consider a counterfactual world, w . At w , there are billions of ^{265}Sg particles, and none of them ever decays. Many of these particles are used to perform 'coin tosses', and since there is never any decay, all of the tosses land Heads. Consider one particular toss, which took place at time t . What was the chance, at t , that the toss would land Heads?

From our present point of view, the answer is: $1/2$. For we know that the particle

involved in the relevant toss is a particle of ^{265}Sg , and we are working on the assumption that part of what it is to be a ^{265}Sg is to have a half-life of 8.9 seconds. The result is that, unlike Lewis, we are able to make room for a possibility whereby fair coins always land Heads. This is because our specification of the relevant world was, in part, a specification of the relevant chances. Since we are assuming that part of what it is to be a particle of ^{265}Sg is to have a half life of 8.9 seconds, when we stipulated that w was to contain particles of ^{265}Sg we were thereby stipulating that the ‘coin tosses’ of w would be fair. And because of our anti-reductionism—because there was no effort to reduce objective chances to further, more fundamental properties—our further stipulation that no ^{265}Sg particle ever decays did nothing to undermine the initial stipulation that coin tosses were to be fair.

A friend of Lewis’s best-system analysis, in contrast, would have faced just such undermining. For once she stipulates that no ^{265}Sg particle ever decays, the best-system for w can be expected to deliver the result that ^{265}Sg is not radio-active (and therefore has an infinite half-life). So she will have undermined the initial stipulation that the relevant particles were of ^{265}Sg (and therefore of a kind that essentially has a half-life of 8.9 seconds).

So much for my example. I hope it makes clear that anti-reductionism is compatible with a non-mysterious epistemology of objective chance, and that anti-reductionism is able to escape the peculiar sort of undermining that burdens Lewis’s best-system analysis.

2 Evidential Support

Let us now turn to the notion of evidential support. The basic idea is straightforward: for evidence E to *support* proposition A is simply for E to constitute better evidence for A than for $\neg A$. More generally, one can think of evidential support as coming in *degrees*: if E necessitates A , then E supports A to degree 1; if E necessitates $\neg A$, then

E supports A to degree 0; for E to support A to an intermediate degree x is for x to be the extent to which E constitutes better evidence for A than $\neg A$, normalized to fall within the $[0, 1]$ interval.

As I will be understanding it here, the notion of evidential support is meant to be *objective* in three different respects. Firstly, it is meant to have *ways for the world to be* as its relata, rather than *ways for the world to be under a mode of presentation*. Suppose, for example, that E supports to degree x the proposition that Hesperus is visible. Then E must also support to degree x the proposition that Phosphorus is visible. This is because the *way the world must be* in order for Hesperus to be visible is precisely the *way the world must be* in order for Phosphorus to be visible: in both cases the world must be such that Venus—the planet itself—is visible.

Secondly, the notion of evidential support is meant to be objective in the sense of being independent of one’s views about the way the world is. Suppose, for example, that an experimental blunder has led us to the mistaken conclusion that ^{265}Sg has a half-life of 4.9 seconds (rather than the correct value of 8.9 seconds). It might then be true that—from our own mistaken perspective—the evidence that a standard (8.9 second) ‘coin toss’ has been set up supports the conclusion that the coin will land Heads to degree $1/4$. It will nonetheless remain the case that—in the objective sense of evidential support that we will be conceded with here—the evidence supports the conclusion that the coin will land Heads to degree $1/2$. Our relation of evidential support is therefore similar to the relation of *confirmation* that was studied by Hempel (1945*a,b*) and Carnap (1962, 1971).⁴

⁴An important difference between the present account and Carnap’s is that, unlike Carnap, we will not be presupposing that the notion of evidential support is constrained solely on the basis of formal considerations. Carnap was interested in finding principles of ‘inductive logic’ that would help reduce the range of rational initial credence functions on the basis of purely formal ‘symmetry’ constraints. He wasn’t however, confident that he would be able to identify enough constraints to single out a unique initial credence function: “We do not know today whether in this future development [of inductive logic] the number of admissible [initial credence functions] will always remain infinite or will become finite and possibly even be reduced to one. Therefore, at the present time I do not assert that there is only one rational [initial credence function].” (Carnap 1971) The notion of inductive probability in (Skyrms 1986) is perhaps closer to the notion of evidential support that we will be discussing here. My own thinking about these issues has been influenced by White (2005, 2010).

Thirdly, the notion of evidential support is meant to be objective in the sense of being independent of one's *cognitive limitations*. Suppose, for example, that we learn that a computer has written down the billionth digit of π (in base 10).⁵ In fact, the billionth digit of π is the number 9. So our evidence *necessitates* the conclusion that the computer has written down the number 9. Accordingly, if E is our evidence and A is the proposition that the computer has written down the number 9, E must support A to degree 1. But now suppose that we are not, as a matter of fact, in a position to carry out the requisite calculation. Then although optimal use of E would mandate assigning a probability of 1 to A , optimal use of E is not within our reach. Perhaps the best that *we* are in a position to do with the evidence, from our limited computational perspective, is assign a probability of 0.1 to A (on the grounds that there 10 different digits in base 10, and that we are not in a position to choose between them on the basis of E). I will not take a stand on the issue of whether it is possible to construct a notion of subjective evidential support which captures this limited perspective. All I wish to emphasize here is that the objective notion of evidential support we will be considering here is not intended to capture limited perspectives of this kind.

One reason to focus on an objective notion of evidential support is that it puts us in a position to think of objective chance as a *special case* of evidential support. More specifically, it puts us in a position to accept the following principle:

CHANCE AS EVIDENTIAL SUPPORT

For the objective chance of A at t to be x just is for *perfect* evidence at t to support A to degree x .

where perfect evidence at t is full information about how the world is up to time t , except for any information about how chance events that occur at or after t turn out. (I will have more to say about the connection between objective chance and evidential support

⁵Here I have greatly benefited from discussion with Adam Elga.

in section 4.)

The preceding remarks are meant to serve as an unofficial introduction to the notion of evidential support. The official story is that the notion of evidential support is to be characterized on the basis of its theoretical role, as given by the following constraints:

1. Framework Constraints:

Our first two framework constraints are meant to ensure that the notion of evidential support has as its relata *ways for the world to be*, rather than *ways for the world to be under a particular mode of presentation*, and that the notion of evidential support is not sensitive to the cognitive limitations of particular subjects.

(a) RELATA

The relation of evidential support is to hold between a proposition E , a proposition A and a real number $x \in [0, 1]$.

(b) PROPOSITIONS

Propositions, in the present context, are to be thought of as sets of metaphysically possible worlds.

Two additional Framework Constraints are meant to ensure that whenever E necessitates A , E supports A to degree 1, and, more generally, that the notion of evidential support behaves like a countably additive probability function, defined over the space of metaphysically possible worlds.

(c) NECESSARY PROPOSITIONS

If \top is the necessary proposition, then E supports \top degree 1.

(d) COUNTABLE ADDITIVITY

Let I be a countable set, and suppose that A_i and A_j ($i, j \in I$) can only be true in the same possible world if $i = j$. Then the degree to which E supports

the union of the A_i ($i \in I$) is the sum of the degrees to which E supports each A_i .

2. Additional Constraints:

Two further constraints constitute the meat of the proposal, and will receive extended discussion below.

(e) LOCALISM

We will refrain from assuming that the notion of evidential support is well-defined for arbitrary E and A , or that it is well-defined in the absence of substantial background assumptions.

(f) THE EVIDENTIAL PRINCIPLE (informal version)

One should align one's credences with one's views about what the evidence supports.

(A more precise statement of this constraint will be supplied in section 4.)

This concludes our official characterization of the notion of evidential support.

Do our constraints supply a *reduction* of the relation of objective chance to 'more fundamental' properties, or a reduction of the relation of evidential support to 'more fundamental' properties? It is not clear to me that it does. For it is not clear to me that one should think of the notion of evidential support as 'more fundamental' than the notion of objective chance, and it is not clear to me that the terms occurring in our theoretical constraints for evidential support are 'more fundamental' than the notion of evidential support. (It is not obvious to me, for example, that the normative notion that occurs in the Evidential Principle should be thought of as 'more fundamental' than the notion of evidential support.)

Fortunately, we need not worry about such issues from our anti-reductionist standpoint. We set ourselves the task of elucidating the notions of objective chance and

evidential support, but we are under no pressure to do so in a way that delivers a reduction.

The basic proposal is now in place. Next on the agenda is to have a closer look at Localism and the Evidential Principle.

3 Localism

We begin with Localism. In section 3.1 I will argue that there is no good reason for thinking that the notion of evidential support would be generally well-defined from a *blank-slate perspective*—a Cartesian point of view, in which one makes no assumptions that could intelligibly be called into question. This will help set the stage for a discussion of Localism in sections 3.2–3.4.

3.1 The Blank-Slate

In this subsection we will consider three examples. Each of them is intended to illustrate a respect in which it is difficult for the notion of evidential support to gain traction from a blank-slate perspective.

Example 1: The Urn

Let E be the proposition that: (i) there is an urn containing balls, (ii) every ball in the urn is either red or blue, and (iii) a ball was drawn from the urn. Let R be the proposition that the ball that was drawn is red. To what degree does E support R , from a blank-slate perspective? There is some temptation to appeal to considerations of symmetry. One might answer ‘to degree 1/2’, on the grounds that there are only two hypotheses—red and blue—and that our blank-slate perspective affords us ‘no reason’ to prefer one of these hypotheses over the other.

To see why this is problematic, consider a slight variation of the case. Everything

is as before, except that this time we learn that the urn contains balls of *three* different colors: red, light blue, and dark blue. To what degree does our revised evidence E' support R —the proposition that a red ball was drawn—from a blank-slate perspective? If we again appeal to considerations of symmetry, it would appear that we should answer ‘to degree $1/3$ ’, on the grounds that we have three hypotheses—red, light blue and dark blue—and that, as before, our blank-slate perspective affords us ‘no reason’ to prefer one of these hypotheses over the rest.

But wait! It is *analytic* (or so I shall assume) that the things that are blue are precisely the things that are either light blue or dark blue. So the evidence in both cases is exactly the same: there is no difference between *the way the world must be* in order for E to be the case and *the way the world must be* in order for E' to be the case. Since evidential support relates *ways for the world to be* rather than *ways for the world to be under a mode of presentation*, this means that the degree to which E supports R must equal the degree to which E' supports R .

It is not clear what grounds one could use to make progress on this issue, from a blank-slate perspective. Any judgment about the ‘symmetry’ of a set of hypotheses relative to one’s evidence presupposes a family of *natural categories*—a family of categories that in some sense constitute an ‘adequate basis’ for epistemological inquiry—and it is not clear how one could go about selecting such a basis in a principled way, from a blank-slate perspective.

Titelbaum (2010, 2011) makes a formal version of this point. He shows that Carnap’s (1962) effort to constrain epistemic support by appeal to purely formal symmetry considerations is inevitably dependent on the particular selection of predicates one chooses to work with. As a result, Carnap’s symmetry constraints can’t really gain traction unless one is able to identify a family of ‘natural predicates’ on which to base one’s theory. And, again, it is not clear how one could go about selecting such a basis in a principled way, from a blank-slate perspective.

Example 2: God's Number

God has selected a positive integer (Arntzenius et al. 2004), using one of the following methods:

The Thin Method

God carries out an ω -sequence of coin-tosses. If the n th toss lands Heads and every other toss lands Tails, God selects number n .

The Thick Method

Consider the result of bending the $[0, 1)$ interval into a loop. By using the Axiom of Choice, God is in a position to divide the loop into countably many regions, any two of which are rotations of one another by some rational distance (Vitali 1905). God maps the resulting regions to the positive integers. God then constructs an ω -sequence of zeroes and ones by carrying out an ω -sequence of coin-tosses, and identifies the real number r within $[0, 1)$ whose binary notation is '0.' followed by the relevant sequence of zeroes and ones.⁶ Finally, God selects the positive integer mapped to the region containing r .

Take either of these two methods, and let E be the proposition that God selected a positive integer using that method. Let A_k be the proposition that the selected number is k . To what degree does E support A_k , from a blank-slate perspective?

Could we say that E supports A_k to different degrees for different k s? Suppose, for example, that we claim that E supports A_k to degree $1/2^k$. That would certainly deliver a probability distribution. But there are many different ways of distributing probabilities unevenly over the natural numbers, and it is not clear that we could have any principled grounds for preferring one over the rest, from a blank-slate perspective. For there is

⁶Some numbers within $[0, 1)$ have two different binary notations (for instance: $1/2$ can be expressed either as '0.1' or as '0.0111...'). If one sees this as introducing a problematic asymmetry, one can simply stipulate that God discards any ω -sequence of coin-tosses that results in a sequence that ends in an uninterrupted sequence of ones.

nothing in either method to suggest that God was more likely to select some integers than others, and a blank-slate perspective would seem to bar the introduction of any additional considerations.

Would a *uniform* distribution do any better?⁷ Could there be an $x \in [0, 1]$ such that, for every k , E supports A_k to degree x ? I had earlier suggested Countable Additivity as one of our Framework Constraints on the notion of evidential support, and Countable Additivity entails that the answer is ‘no’.⁸ But what if we revised this assumption? What would happen if we chose to work with a notion of evidential support that was only *finitely* additive? Doing so is mathematically possible, but, as we shall see, it comes at a cost (Seidenfeld & Schervish 1983, Schervish et al. 1999).

A first reason failures of Countable Additivity come at a cost is that they restrict the range of applications of the notion of epistemic support. In particular, they restrict one’s ability to base a *decision theory* on the notion of epistemic support. For whenever a probability distribution fails to be countably additive, it is possible to construct a ‘robust’ Dutch Book for a subject whose credences are given by that distribution (Williamson 1999, Easwaran 2013).⁹ So a notion of evidential support that failed to satisfy Countable Additivity would lead to the unpleasant result that a subject whose credences are per-

⁷It is worth noting that it is not obvious that we have principled reasons to prefer uniformity from a blank-slate perspective, since such a preference would seem to presuppose that the family of properties *is identical to k* (for k a positive integer) are natural enough to constitute an ‘adequate basis’ for epistemological inquiry. And, as emphasized in the previous section, it is not obvious that one would be entitled to such judgements from a blank-slate perspective.

⁸*Proof:* If $x = 0$, it follows from Countable Additivity that the union of the A_k is supported by E to degree 0, which contradicts the fact that A is *entailed* by E and should therefore be supported to degree 1; if $x > 0$, it follows from Countable Additivity that E supports to the union of the A_k to an infinite degree, which contradicts the Framework Constraints of section 2.

⁹More specifically, if $P(X)$ is the relevant probability distribution, there is a partition \mathcal{E} and a pair $\langle \Theta_i, \epsilon_i \rangle$ for each E_i in \mathcal{E} (where $\epsilon_i > 0$ if $\Theta_i \neq 0$ and $\epsilon_i = 0$ if $\Theta_i = 0$), such that the subject would suffer a sure loss by purchasing every bet B_i (which costs $\$P(E_i)\Theta_i - \ϵ_i , and pays $\$\Theta_i$ if E_i and nothing otherwise), even though the purchase of B_i has a *positive* expected value for the subject whenever $\Theta_i \neq 0$. (When $\Theta_i = 0$, B_i isn’t a substantive bet because it costs nothing and it pays nothing). In order for a Dutch Book of this kind to count as *robust* we impose the additional requirement that the purchasing of every B_i does not require an infinite amount of money to change hands. (As noted in (McGee 1999), a failure to satisfy the additional condition would trivialize the problem by allowing for Dutch Books regardless of the additivity constraints endorsed by the subject.)

fectly aligned with what the evidence supports is nonetheless robustly Dutch-Bookable. In contrast, it is *impossible* to build a robust Dutch Book for a subject whose credence distribution fails to satisfy Uncountable Additivity but satisfies Countable Additivity, along with the standard probability axioms (Easwaran 2013).

What if we took a step further, and revised the assumption that degrees of support are real numbers? Perhaps we could claim that for each k , E supports A_k to an *infinitesimal* degree ι . As it turns out, infinitesimals won't solve our problem: a subject who assigns credence ι to each A_k can still be robustly Dutch-Booked (Arntzenius et al. 2004).

We have seen that abandoning Countable Additivity restricts one's ability to use probability distributions for decision-theoretic purposes. But probability distributions that lack Countable Additivity also have some awkward *formal* properties. Here is an example, which I borrow from Easwaran (forthcoming). (Readers uninterested in the details are welcome to skip this paragraph and the next.) The following probability distribution over the set \mathbb{Z}^+ of positive integers is as sensible as any I can think of:¹⁰

$$P(X|Y) =_{df} \lim_{n \rightarrow \infty} \frac{|X \cap \{1, 2, \dots, n\}|}{|Y \cap \{1, 2, \dots, n\}|}$$

$$P(X) =_{df} P(X|\mathbb{Z}^+)$$

$P(X)$ is finitely additive, but fails to respect Countable Additivity because $P(\mathbb{Z}^+) = 1$ even though $P(\{k\}) = 0$ for each $k \in \mathbb{Z}^+$. It also has the following awkward property: there is a set S and a partition E_i of \mathbb{Z}^+ such that $P(S) = 0$ even though $P(S|E_i) = 1/2$ for each E_i . (Let S be the set of perfect squares, and for each $i \in \mathbb{Z}^+$ such that i is not a power of any other positive integer, let E_i be the set of powers of i .)

¹⁰This definition privileges the 'natural' ordering of the positive integers. One's choice of ordering is significant because one can get different results by working with different orderings. It is not obvious, however, that one could find a principled basis for choosing any particular ordering from a blank-slate perspective. So it is not obvious that a friend of the blank-slate perspective could find a principled basis for working with $P(X|Y)$, rather than a distribution based on a different ordering. I discuss the example nonetheless because it illustrates a general problem for probability distributions that fail to be countably additive.

As it turns out, awkward results of this kind are inevitable: if a probability distribution on a countable space fails to satisfy Countable Additivity it is *impossible* to extend it to a conditional probability distribution that respects Conglomerability: the principle that if the A_j form a partition, then $\min[P(X|A_j)] \leq P(X) \leq \max[P(X|A_j)]$ (Schervish et al. 1984, Hill & Lane 1985). In contrast, there is significant room to maneuver when it comes to probability distributions that fail to satisfy Uncountable Additivity but satisfy Countable Additivity (Kolmogorov 1960, Easwaran forthcoming).

The lesson of all this is that although one can extend the range of cases in which probability distributions are well-defined by abandoning Countable Additivity, one ends up making the notion of a probability distribution less interesting in the process. So although it is certainly possible to give up Countable Additivity, doing carries the risk of a Pyrrhic victory.

We have been considering the question of whether the notion of evidential support is well-defined from a blank-slate perspective, when our evidence consists of the proposition that God has selected a positive integer using one of the methods described above. Our conclusion is that a positive answer to this question would come at a cost. One must either assign different degrees of evidential support to different A_k —and thereby introduce an asymmetry where a blank-slate perspective warrants none—or water-down the notion of evidential support by abandoning Countable Additivity. This gives us some reason to embrace the conclusion that the notion of evidential support is simply not well-defined in the case of God’s Number.

Example 3: The Intuitionist and the Classicist

Consider a debate between a classical logician and an intuitionist. The classicist thinks that for it to be the case that $\neg\neg p$ *just is* for it to be the case that p . The intuitionist disagrees. She thinks that in satisfying the condition that $\neg\neg p$ the world does not *thereby* satisfy the condition that p , because the condition that p imposes more stringent demands

on the world than the condition that $\neg\neg p$.

Each of our logicians sees the other as deploying the wrong categories in constructing her conception of the way the world is. From the classicist's point of view, the intuitionist is making *a distinction without a difference*, by distinguishing between its being the case that $\neg\neg p$ and its being the case that p . From the intuitionist's point of view, the classicist *conflates* distinct conditions, by failing to distinguish between the condition that $\neg\neg p$ from the more stringent condition that p .

This means that there is an important sense in which is impossible to use the categories endorsed by one of our logicians to capture the point of view of the other. Consider, first, the question of how the classicist might model the intuitionist's claim that the condition that p imposes more stringent demands on the world than the condition that $\neg\neg p$. The classicist cannot account for the purported space between the condition that p and the condition that $\neg\neg p$ by introducing 'scenarios' in which it is the case that $\neg\neg p$ without being the case that p . For a scenario in which it is not the case that p is, from the perspective of the classicist, a scenario in which it is the case that $\neg p$. So the classicist would be modeling the intuitionist's position by introducing a scenario that she takes to satisfy both the condition that $\neg p$ and the condition that $\neg\neg p$. But, by the classicist's lights, such a scenario is *absurd*. So, by her own lights, she will have failed to do justice to the intuitionist's point of view.

Nor, incidentally, would it do to model the intuitionist's point of view in *cognitive* terms. In particular, the classicist would fail to do justice to the intuitionist's position if she modeled it as the claim that it is possible for a subject to be in a *cognitive state* whereby she accepts the claim that $\neg\neg p$ without accepting the claim that p . For the classicist might well *agree* with the intuitionist that the transition from $\neg\neg p$ to p requires a non-trivial cognitive accomplishment. That is *not* what the disagreement between our two logicians is about.

What the classicist would have to do in order to truly do justice to the intuitionist's

position to is refrain from making the assumption that for the world to satisfy the condition that $\neg\neg p$ is *already* for it to satisfy the condition that p . But in refraining from making such an assumption she will have given up a crucial part of classical logic. So whereas one can model rival views about, say, astronomy without thereby giving up one's current astronomical views, it is not obvious that one can genuinely do justice to rival views about logic without thereby abandoning one's own logic.

Now consider the question of how the intuitionist might model the classicist's claim that for it to be the case that $\neg\neg p$ is *already* for it to be the case that p . It won't do for the intuitionist to model such a view by taking the classicist to accept the claim that $\neg\neg p \supset p$ (and thereby rule out as *unactualized* a scenario in which it is the case that $\neg\neg p$ without being the case that p). For the intuitionist might well *agree* with the classicist that no such scenario is actualized. That is *not* what the disagreement between our two logicians is about.

What the intuitionist would have to do in order to truly do justice to the classicist's position is refrain from making the assumption that the condition that p imposes more stringent demands on the world than the condition that $\neg\neg p \supset p$. But in refraining from making such an assumption she will have given up an important part of intuitionism. So, again, unlike the case of astronomy, it is not obvious that one can genuinely do justice to rival views about logic without thereby abandoning one's own logic.

I have been arguing that there is an interesting sense in which the adoption of one of our logics leaves no conceptual space for the other, by forcing one to adopt categories that make it impossible to do justice to the rival perspective. With this as our background, let us return to our discussion of evidential support from a blank-slate perspective.

I earlier defined the blank-slate perspective as a Cartesian point of view, whereby we refrain from assuming anything that could intelligibly be called into question. Let us now consider the question of what to say about *logic*, from a blank-slate perspective. It is tempting to think that the issue is straightforward. After all, the philosophical

literature is filled with interesting disagreement about which logic to adopt (Dummett 1991, Wright 2001, Priest 2006, Field 2008). Doesn't this imply that one's own choice of logic is something that can be intelligibly called into question, and therefore something that cannot be taken for granted from a blank-slate perspective?

The lesson of our discussion of classicism and intuitionism is that the issue is not as straightforward as one might have hoped. There is certainly a sense of intelligibility whereby one can treat a rival logic as an intelligible position: a journal-referee might conclude, for example, that an article defending a logic other than her own is 'intelligible' enough to count as a respectable addition to the philosophical literature. But there is also a sense of intelligibility whereby rival logics are *not* intelligible positions. We have seen that the adoption of a logic can entail the adoption of a system of categories that leaves no conceptual space for rival perspectives. We have seen, in particular, that the classicist is committed to treating any space between $\neg\neg p$ and p as *absurd*. So there is a sense of intelligibility whereby the classicist is committed to treating intuitionism as 'unintelligible' (because absurd), even if she thinks that a paper defending intuitionism could be a respectable contribution to the literature.

Our two different senses of intelligibility give rise to two different senses of 'blank-slate perspective'. Let us say that from an *absolute* blank-slate perspective one must refrain from assuming anything that could be intelligibly called in to question in the weaker sense of intelligibility, whereby a respectable contribution to the philosophical literature would count as intelligible. In contrast, let us say that from a *moderate* blank-slate perspective it is only necessary to refrain from assuming anything that could be intelligibly called in to the question in the stronger sense of intelligibility, whereby the classical logician would count the intuitionist's position as unintelligible.

It is hard to see how one could presuppose much by way of of a logic from an absolute blank-slate perspective. For just about any principle of reasoning one can think of has been called into question by a respectable contribution to the philosophical literature.

And without principles of reasoning to work with, it is not clear how a notion of epistemic support could be well-defined. We simply don't have enough constraints in place to gain traction on the project of using our evidence to make principled discriminations amongst ways for the world to be. (With no principles of reasoning to work with, even Descartes's *cogito* argument would seem to be out of reach.)

A moderate blank-slate perspective, in contrast, allows for principles of reasoning that are not called into question. The classicist, for example, will not see intuitionism as an obstacle to presupposing classical logic from a moderate blank-slate perspective. One can therefore hope to gain some traction on questions of epistemic support. But such progress comes at a cost. For in taking principles of reasoning for granted, one ceases to operate from *fully neutral* perspective.

Notice, in particular, that in treating intuitionism as absurd (and therefore as an illegitimate ground for casting doubt on classical principles of reasoning), the classicist doesn't just presuppose that intuitionism is mistaken. She presupposes that it is *beyond consideration*. More specifically: since the classicist thinks that classical logic can never take one from non-absurd premises to an absurd conclusion, and since she thinks that intuitionism would be an absurd conclusion, she must think that—regardless of the evidence one is presented with—classical logic will never lead to the conclusion that intuitionism is correct. (This is not to say that it is impossible to start from a classicist perspective and end up an intuitionist. But it is to say that it is impossible to do so without thereby abandoning the very constraints that allowed the classicist to gain traction on the question of how to use one's evidence to make principled discriminations between ways for the world to be.)

It is initially tempting to think that the notion of evidential support could gain traction from a fully neutral perspective. But we have seen that an *absolute* blank-slate perspective is too unconstrained to allow for facts about evidential support to be well-defined. One is trapped between a rock and a hard place: one must choose between

operating on the basis of non-neutral assumptions and having too few constraints for the notion of evidential support to gain traction.

Lessons Learned

The examples we have been discussing have a common theme. In all three cases, I argued that a workable notion of epistemic support comes at a cost. For the very constraints that make questions of evidential support worth investigating place limits on what the notion of evidential support can be used to achieve.

In the case of the urn, we found that we had to choose between presupposing a set of ‘natural categories’ (and thereby limiting the range of application of the notion of evidential support by being dogmatic about what counts as a natural category), and operating with so few constraints that it is hard to conclude anything of interest on the basis of our evidence. In the case of God’s number, we found that we had to choose between working with Countable Additivity (and thereby limiting the range of cases with respect to which probability distributions are well-defined), and operating with a watered-down version of probability theory, which is less interesting as a theoretical tool. In the case of the rival logicians, we found that we had to choose between taking certain logical principles for granted (and thereby limiting the range of application of the notion of evidential support by being dogmatic about our logical principles), and operating with so few constraints that it is hard for to gain any traction on questions of evidential support.

3.2 Epistemology Games

We have seen that the project of using one’s evidence to make principled discriminations amongst ways for the world to be requires a choice between *suboptimal alternatives*: the very constraints that make the project interesting must also limit its scope.

Deciding how to constrain the notion of epistemic support is in some ways analogous

to deciding which game to play. Notice, in particular, that in deciding which game to play, one is also forced to choose between suboptimal alternatives, since the very rules that make a game desirable in some ways, can make it undesirable in others. The complexity of game-play that results from the rules of chess, for example, is precisely what makes the game appealing in some contexts, but it is also what makes it unappealing in others. As a result, it would be a mistake to think that there is a context-independent answer to the question of what game to play. It all depends on one's aims, and on the situation at hand. (Is one's objective to show off, or is to get everyone to have a good time? Does one have access to a chessboard and a willing opponent?)

Similarly, there is room for thinking that there is no context-independent answer to the question of how to constrain the notion of evidential support. As in the case of deciding which game to play, it may depend on one's aims, and on the situation at hand. Suppose, for example, that one's aim is to find principled reasons for deciding whether to adopt classical logic or intuitionism. As we have seen, an absolute blank-slate perspective is unworkable: unless one is prepared to take principles of reasoning for granted, one can't even get started on the project of making principled discriminations amongst ways for the world to be on the basis of one's evidence. Taking *classical logic* for granted would have the advantage of delivering real traction to questions of epistemic support. But it would come at a huge a cost in the relevant context. For, as we have seen, a classical perspective would commit us to counting intuitionism as absurd, and would therefore take intuitionism off the table from the start.¹¹

Now consider a different sort of project. Suppose that one is setting up a casino, and that one is hoping to design a system of odds that will entice one's customers while delivering the result that, with very high probability, the casino will end make money in

¹¹I don't mean to suggest that working within a classical would be useless for someone trying to decide between classical logic and intuitionism. One might, for example, do one's best to address a philosophical paradox from the perspective of classical logic, and turn to intuitionism in frustration, only to find that the paradox is absent from an intuitionistic perspective. The point is simply that intuitionism could never be the outcome of an investigation *within classical logic*.

the long run. In this case, one might find that the cost of working with a classical system is not too high a price to pay. I do not claim that there is no cost at all: there may well be a proof in some non-classical system to the effect that assigning odds based on classical logic is likely to lead to disastrous losses. (The casino owner might ask, “What if such a non-classical logic turns out to be right?”, and thereby see herself as taking some sort of risk.) But the cost of remaining open-minded about which logic to accept would presumably be much higher. For unless she is willing to commit to some logical system or other—even if the commitment is provisional—she will lack the theoretical tools to get started on the project of identifying an adequate system of odds.

If our casino owner is to make any progress, she will also need to make some *material assumptions*. Unless she assumes, for example, that properties corresponding to card-suits and card-numbers can be counted ‘natural categories’, in the sense described above, it will be hard for her to come up with a principled assignment of initial probabilities. It would also be hard to come up with a principled assignment of probabilities if one treated certain *non-standard scenarios* as live options (for instance, a scenario in which one of the cards has been misprinted in a way that makes it ambiguous), or if one treated *skeptical scenarios* as live options (for instance, a scenario in which outcomes are manipulated by an evil demon). Rather than face the epistemological paralysis that would result from treating such scenarios as live options, the casino owner might decide to rule them out from the start—even though doing so comes at a cost, since it is certainly possible for one of them to be actualized, and the casino owner will be taking a risk by failing to take them into account in her calculations.

Let an *epistemology game* be a set of constraints on the notion of evidential support: a set of background assumptions with respect to which the project of using one’s evidence to make principled discriminations amongst ways for the world to be can be expected to gain traction. An epistemology game can include ‘framework constraints’—i.e. constraints aimed at shaping the general behavior of the notion of evidential support—such as a

choice of logic. It can also include ‘material constraints’, such as assumptions about the properties that are to be counted as ‘natural categories’, or assumptions to the effect that certain non-standard outcomes and skeptical scenarios are to be barred from consideration.

The notion of an epistemology game is, of course, indebted to Wittgenstein’s *On Certainty*, since one can think of the adoption of an epistemology game as the adoption of ‘hinge propositions’ on which to base one’s epistemological inquiry. My own views have been greatly influenced by (Wright 2004).¹²

Let us take stock. The upshot of subsection 3.1 can now be stated as follows: it is hard for the notion of epistemic support to gain traction unless a suitable epistemology game is in place. This leaves us with the question of what epistemology game to adopt. What I suggested in the present section is that there may be no context-independent answer to this question. It may all depend on one’s aims, and on the situation at hand. We will next consider the question of whether there is such a thing as an ‘objectively correct’ epistemology game—or, at least, such a thing as an epistemology game which is ‘objectively correct’ relative to a given context.

3.3 Objectivity

One might think that classical logic is objectively correct, and therefore that an epistemology game can only be objectively correct if it is based on classical modes of reasoning. One might go further, and take the view that the Framework Constraints of section 2 are objectively correct. One might go further still, and take the view that there are objectively correct *material* constraints: one might think, for example, that there is a class of properties that constitute objectively natural categories, and therefore an objectively

¹²According to Wright, however, there are circumstances under which one is ‘entitled’ to take for granted assumptions that are needed for inquiry to gain traction. Here I depart from Wright by remaining neutral on the question of whether one might, in some sense, be entitled to the adoption of a given epistemology game.

correct basis for epistemic inquiry, in the sense of section 3.1.

Although I would be delighted if there turned out to be a notion of objective correctness that could support such claims, my own efforts to spell out such a notion have been unsuccessful (Rayo 2013, ch. 2). The good news is that there is no need to take a stand on this issue for present purposes. Throughout this paper, I will remain *neutral* on the issue of whether it makes good sense to describe an epistemology game as ‘objectively correct’. I would, however, like to say a few words about how to think about epistemology games in the absence of an objective notion of correctness.

It is useful to start with an analogy. Consider the rule: *you lose the game if your king is captured*. Everyone can agree that this is correct as a rule of chess. But it would be odd to ask whether the rule is correct in a *transcendent* sense, a sense that does not depend on which game is being played. For it is natural to think that rather than having a transcendent notion of correctness for games, what we have is a variety of ‘local’ notions of correctness, one for each game. It is by setting forth a board game that one generates the conditions that make it possible to make sense of a distinction between correct and incorrect rules.

Similarly, there is room for thinking that rather than having a transcendent notion of correctness for epistemology games, what we have is a variety of ‘local’ notions of correctness, one for each epistemology game. For there is room for thinking that it is by setting forth an epistemology game that one generates the conditions that make it possible to make sense of a distinction between correct and incorrect ways using one’s evidence to make principled discriminations amongst ways for the world to be.

Even if there turns out to be no transcendent notion of correctness, however, it would be a mistake to conclude that one’s decision about what what epistemology game to adopt need not be sensitive to the way the world is. This is because what sorts of epistemology games would serve one’s purposes depends, in part, on the way the world is. In the absence of a transcendent notion of correctness, the question of what epistemology

game to adopt cannot be fully separated from pragmatic considerations, in much the way that the question of what board game to play cannot be fully separated from pragmatic considerations.

3.4 Back to the Constraints

In section 2 I set forth two sets of constraints on the notion of epistemic support: ‘Framework Constraints’ and ‘Additional Constraints’. Before bringing our discussion of localism to a close, I would like to say a few words about how these constraints interact with our present discussion.

The ‘Framework Constraints’ are best thought of as specifying the general outline of an epistemology game. They shape the notion of evidential support by ensuring that it is based on probabilistic principles of reasoning (which I tacitly assume to be classical), and ensuring that evidential support applies to *ways for the world to be*, rather than *ways for the world to be under a mode of presentation*.

In describing the Framework Constraints as partial specifications of an epistemology game I wish to emphasize three points:

1. Non-Neutrality

I argued in section 3.1 that in taking classical logic for granted one fails to take a fully neutral perspective. For similar reasons, one fails to take a fully neutral perspective when one works against the background of the (classical) probabilistic framework that is embodied in our Framework Constraints.

This lack of neutrality comes at a cost. It means that certain hypotheses are bared from consideration in spite of being intelligible, in the weak sense of intelligibility discussed above. It ensures, for example, that one’s epistemology game will never deliver the conclusion that one should embrace intuitionistic logic, or a non-standard probability theory. I would like to suggest, however, that—in a large range

of circumstances, and relative to a broad family of aims—this will turn out to be a price worth paying. For our Framework Constraints deliver significant traction on the notion of evidential support in exchange for a loss in neutrality that will often seem relatively minor.

2. Context-Sensitivity

I do *not* claim that the Framework Constraints deliver a useful way of constraining the notion of evidential support every context. We have already discussed one case in which they don't—the case in which one is engaged in the project of deciding which logic to accept—but there are many other examples. There are, for instance, cases in which one is simply not in a position to come up with a workable assignment of probabilities, or in which one is in a position to do so but in which doing so would require more effort than would be warranted by the situation at hand.

3. Incompleteness

Even in contexts in which the Framework Constraints prove to be useful, they are unlikely to deliver a workable epistemology game on their own. In most cases, one will also need 'material constraints' of the kind I described above. The Framework Constraints are therefore best thought of as supplying the initial scaffolding on which a full-fledged epistemology game might be built.

What about the Localism and the Evidential Principle? Although the issue is partly terminological, I think it is not particularly useful to think of these principles as helping set up an epistemology game. Localism is best thought of as part of the story of what epistemology games consist in, rather than as part of the specification of any particular epistemology game. And the Evidential Principle is so central to our understanding of epistemic support that I think the result of giving it up wouldn't be to switch epistemology games; it would be to go off on a different topic.

4 The Evidential Principle

In this section we will discuss the Evidential Principle, which is closely related to Christensen’s ‘Rational Reflection Principle’ (Christensen 2010).¹³

It will be useful to start with some notation. I will use ‘ $sup_G(A|E) = x$ ’ as a formalization of the claim that evidence E supports proposition A to degree x relative to epistemology game G , and will hereafter restrict my attention to the special case in which G satisfies the Framework Constraints of section 2.

Consider a subject who adopts epistemology game G ; in other words: she takes for granted that the principles in G are correct. The Evidential Principle is the claim that such a subject should align her credences with her views about what her evidence supports, by the lights of G . More precisely:

THE EVIDENTIAL PRINCIPLE

$$c_G(A|E \cdot [sup_G(A|E) = x]) = x$$

where A and E are arbitrary propositions, c_G is rational initial credence function for a subject who adopts G , and $[sup_G(A|E) = x]$ is the proposition that E supports A to degree x relative to G .

Informally, one might think of the Evidential Principle as stating that one’s credences should match one’s *expectation* of what the evidence supports.¹⁴

¹³Christensen’s principle is $c(A|[Pr(A) = n]) = n$, where c represents an agent’s credences and Pr represents the credences that would be maximally rational for someone in that agent’s epistemic situation. Christensen uses the principle to motivate a puzzle that can be seen as bringing out a tension between what it is rational to believe, and what it is rational to believe that it is rational to believe. It is not clear to me, however, that Christensen’s puzzle would arise if $Pr(A)$ was interpreted as capturing evidential support, in the present sense. (For instance: I see no obvious reason to resist the idea that each of the evidential situations in Christensen’s example supports to degree 1 the true proposition about the position of the clock.)

¹⁴This is because, given a simplifying assumption, the Evidential Principle is equivalent to:

CREDENCE AS EXPECTATION OF EVIDENTIAL SUPPORT

$$c_G(A|E) = \sum_{i \in I} x^i \cdot c_G([sup_G(A|E) = x^i]|E)$$

where each of the $[sup_G(A|E) = x^i]$ ($i \in I$) is a hypothesis about the extent to which her evidence supports of A .

The Evidential Principle imposes constraints on the ways in which one's assignment of credences to propositions about evidential support should interact with one's assignment of credences to other propositions. Suppose, for example, that one assigns credence 0.5 to the proposition that one's evidence supports A to degree 0.6, and credence 0.5 to the proposition that one's evidence supports A to degree 0.8. Then the Evidential Principle entails that one should assign credence $(0.6/2 + 0.8/2) = 0.7$ to A (all of this, of course, relative to whichever epistemology game one adopts, a qualification that will sometimes be left tacit in what follows).

The purpose of this section is to examine The Evidential Principle in greater detail, and to describe its relation to the Principal Principle. I begin with some general remarks.

Rationality

The Evidential Principle delivers a necessary condition on doxastic rationality. It tells us that one's use of evidence can only count as rational (relative to an epistemology game G) to the extent that one succeeds in aligning one's credences with one's expectation of what the evidence supports (relative to G).

Although the Evidential Principle is *neutral* on the question of whether there is an 'objectively correct' epistemology game, a friend of objective correctness might wish to claim that there is a sense of 'ought' in which any subject *ought* to accept the 'objectively

Here are the details. The Evidential Principle entails Credence as Expectation of Evidential Support on the assumption that at most countably many hypotheses $[sup_G(A|E) = x^i]$ ($i \in I$) are compatible with E . For in that case:

$$\begin{aligned} c_G(A|E) &= \sum_{i \in I} c_G(A|E \cdot [sup_G(A|E) = x^i]) \cdot c_G([sup_G(A|E) = x^i]|E) && \text{(by countable additivity)} \\ &= \sum_{i \in I} x^i \cdot c_G([sup_G(A|E) = x^i]|E) && \text{(by the Evidential Principle)} \end{aligned}$$

Credence as Expectation of Evidential Support entails the Evidential Principle on no assumptions. For by substituting $E[sup_G(A|E) = x]$ for E in Credence as Expectation of Evidential Support one gets:

$$c_G(A|E[sup_G(A|E) = x]) = \sum_{i \in I} x^i \cdot c_G([sup_G(A|E) = x^i]|E[sup_G(A|E) = x]) = x.$$

correct' epistemology game. She could then use the Evidential Principle to conclude that—in the relevant sense of subject 'ought'—any subject *ought* to align her credences with her views about what her evidence supports, by the lights of the 'objectively correct' epistemology game. A normative notion of this kind would go beyond the notion of doxastic rationality that we will be discussing here.

Action

The Evidential Principle helps establish a link between the notion of evidential support and the notion of rational action. This is because there is a link between the notion of *credence* and the notion of rational action, and the Evidential Principle helps extend this link to the notion of evidential support.

The precise character of the link between credence and rational action will depend on the details of one's decision theory, but the following might be used as a rough working hypothesis. Start by using the subject's credences (and preferences) to characterize her *expected utilities*. Then set forth the following principle:

THE ACTION PRINCIPLE

For a subject to act rationally *just is* for her actions to maximize her expected utilities.

The Evidential Principle and the Action Principle together entail that one should act in accordance with one's views about what the evidence supports. A little more carefully: the Evidential Principle entails that doxastic rationality requires one to align one's credences—and therefore one's expected utilities—with one's views about what the evidence supports; the Action Principle entails that pragmatic rationality requires one to align one's actions with one's expected utilities. So full rationality requires one to align one's actions with one's views about what the evidence supports.

A Normative Conception of Credence?

Credences are often thought of as degrees of *belief* or degrees of *certainly*. This makes it natural to suppose that in describing a subject's credences one is describing her *psychological state*. It seems to me, however, that this is a potentially dangerous way of thinking about credences.

Formal epistemology has taught us that credences supply a useful framework for theorizing about rational constraints on belief, and decision theory has taught us that they supply a useful framework for theorizing about rational constraints on action. So there is a good case to be made for the claim that talk of credences is useful in *normative* contexts. In contrast, it is not at so clear that talk of credences is in good order as a purely *psychological* hypothesis. For, as Holton (forthcoming) points out, it is not at all clear that our best cognitive science—our best account of how the mind works—should ultimately make room for degrees of certainty.

In treating an ascription of credences as a psychological hypothesis, it seems to me that one would be taking a good normative theory, and making it hostage to empirical fortune. Suppose that Holton is right, and that credence-talk has no place in our best cognitive science. It would be a shame to conclude, on those grounds alone, that credences have no role to play in our *normative* lives. It would better to work with a notion of credence that allows it to play the normative role without being burdened by empirical commitments.

One way of developing this idea is to start with the thought—defended in (Eriksson & Hájek 2007)—that the notion of credence is simply whichever notion “forms the basis of our best theory of rational belief and decision”. On my preferred implementation of the proposal—which differs from Eriksson and Hájek's—one thinks of the notion of credence as *implicitly defined* by the Evidential Principle and the Action Principle.

On the resulting proposal, credences might be thought of as a device for the normative

assessment of subjects. The point of ascribing a credence-function to a subject is not to make a claim about the subject’s psychology. It is to help keep track of what would count as an optimal use of the subject’s evidence, and use the resulting piece of theory to figure out how it would be rational for the subject to act, given a hypothesis about her preferences. The outcome of one’s theory can then be used to assess the normative status of the subject’s actions. For we can say that the subject has acted *as she should have*—both epistemically and practically—to the extent that her actual actions match the outcomes of our normative theorizing. Whether or not the subject is best understood as having a psychologically realized credence-function can then be regarded as a separate matter, to be decided by empirical investigation.

Even though I think that this normative conception of credence deserves to be taken seriously, I will not be presupposing that it is correct in the present context.

4.1 The Principal Principle

I will end my discussion of the Evidential Principle with a discussion of its relationship to the Principal Principle (Lewis 1980).

The Principal Principle is a statement of the relationship between objective chance, rational credence and *admissibility*. What is admissibility? For a proposition to be admissible at t (or for it to be t -admissible) is for it to include no information “about how chance events in the present and future [relative to t] will turn out” (Lewis 1994).

Before we can state the Principal Principle we need some notation. Let A be an arbitrary proposition, let χ_t^i be a probability distribution, let E_t^i be a t -admissible proposition which entails that the objective chance of A is $\chi_t^i(A)$, and let c be a rational initial credence function. Then one can state the Principal Principle as follows:¹⁵

¹⁵So as to simplify the discussion of undermining below, I follow (Lewis 1994) in formulating the Principle so that its application is explicitly restricted to *admissible* information about the objective chances. Lewis’s original (1980) formulation of the principle includes no explicit mention of this constraint, but is an easy consequence of the version of the principle presented here.

THE PRINCIPAL PRINCIPLE

$$c(A|E_t^i) = \chi_t^i(A)$$

As Lewis (1994) points out, the reductionist about objective chances is committed the possibility of an *undermining future*: a future which differs from the actual future because of differences in the ways chance events turn out, and which is such that if it came to pass, it would determine a chance distribution different from the one that is determined by the actual future. (Given suitable background assumptions, an example of an undermining future is a future in which all ^{265}Sg ‘coin tosses’ land Heads from now on—i.e. a future in which ^{265}Sg particles cease to decay.)

The possibility of undermining futures entails that information about the objective chances at a time is not generally admissible at that time. (Notice, in particular, that if the current chances are incompatible with an undermining future F , they thereby give us information about how future chance events will turn out: they tell us that chance events will turn out somehow other than how they turn out in F .)

This is a problem because when information about the objective chances fails to be admissible, the Principal Principle cannot be applied. (More specifically, the Principle only delivers the result that $c(A|E_t^i) = \chi_t^i(A)$ when E_t^i is an *admissible* proposition which carries information about the objective chances.) As a result, the reductionist is seriously limited in her ability to apply the Principal Principle. She is, in effect, committed to the claim that “if chancemaking patterns extend into the future, then any use of the Principal Principle is fallacious” (Lewis 1994).

Lewis’s initial reaction to this problem was to keep the Principal Principle intact, and modify reductionism so as to get the result that information about the objective chances at a time is always admissible at that time (Lewis 1980, 1986*b*).¹⁶ But he later changed

¹⁶This is done by claiming that the history-to-chance conditionals (whose antecedent is a complete specification of historical facts up to a given time and whose consequent is a complete theory of chance at that time) are always *necessarily* true. Since Lewis thinks that a complete specification of historical facts up to a given time is always admissible at that time, this entails that the complete theory of chance

his mind, preferring to retain his preferred brand of reductionism—and therefore the view that information about the objective chances at a time is not generally admissible at that time—and compensate by working with a modified version of the Principal Principle (Lewis 1994, Hall 1994).

An advantage of anti-reductionism is that one need not take this last step. The anti-reductionist faces no pressure to think that information about the objective chances is inadmissible, and is therefore in a position to work with the original version of the Principal Principle.

In section 2 I claimed that objective chance is best thought of as a special case of evidential support. I claimed, in particular, that objective chance is what *perfect* evidence supports, and labeled this claim ‘Chance as Evidential Support’. Our discussion of localism in section 3 brings out the need for an important qualification: to the extent that the notion of evidential support is only able to gain traction with respect to an epistemology game, one should think that the notion of objective chance is only able to gain traction with respect to an epistemology game. So our thesis is best put as follows: the objective chances relative to G are what the perfect evidence supports relative to G .

If our account of evidential support is to do justice to this thesis, however, it is not enough to set forth Chance as Evidential Support as an axiom, and leave it at that. We need to show that Chance as Evidential Support delivers a conception of objective chance that plays the right *theoretical role*. As I noted earlier, Lewis once thought that the Principal Principle “capture[s] all we know about chance”. If he was right—and I believe that he was—we need to show that the Principal Principle is a consequence of Chance as Evidential Support.

Fortunately, the Evidential Principle ensures that this is so. When information about the objective chances at a time is assumed to be admissible at that time, the Principal Principle (relativized to G) follows from the Evidential Principle (relativized to G),

for a given time is always admissible at that time.

together with Chance as Evidential Support. (See appendix for proof.)

5 Conclusion

The aim of this paper has been to develop an account of evidential support.

My account is *anti-reductionist*: it does not presuppose that the notion of evidential support is reducible to ‘more fundamental’ notions. This allows it to escape certain awkward consequences of reductionist positions such as Lewis’s.

My account is *localist*: it does not presuppose that the notion of evidential support is well-defined independently of substantial background assumptions. I argued that the notion of evidential support is usually only well-defined against the background of an *epistemology game*, and suggested that it is not clear that it makes sense to speak of an epistemology game as being ‘objectively correct’.

My account of epistemic support includes an account of objective chance. I argued that objective chance can be seen as a *special case* of evidential support—objective chance is what perfect evidence supports—and buttressed the view by noting that there is a natural way of deriving Lewis’s Principal Principle from basic principles about evidential support.

Appendix

I claimed in section 4.1 that when information about the objective chances at a time is assumed to be admissible at that time, the Evidential Principle (relativized to G) and Chance as Evidential Support together entail the Principal Principle (relativized to G).

When information about the objective chances at a time is allowed to count as admissible at that time, the Principal Principle can be restated as follows:

THE PRINCIPAL PRINCIPLE (Reformulated)

$$c(A|E_t \cdot [\chi_t(A) = x]) = x$$

where E_t is a t -admissible proposition and $[\chi_t(A) = x]$ is the proposition that the objective chance at t of A is x .

We shall therefore show that $c_G(A|E_t \cdot [sup_G(A|E_t) = x]) = x \rightarrow c_G(A|E_t \cdot [\chi_t(A) = x]) = x$, where E_t is t -admissible. We make use of Chance as Evidential Support, and assume that information about the objective chances at t is admissible at t .

Let I be the set of i such that Q_t^i is a hypothesis about the perfect evidence at t that is compatible with E_t . (For technical reasons, I shall assume that I is *countable*; I shall also assume that the Q_t^i are such as to guarantee that they count as perfect evidence at any world in which they obtain.)

Let I_{A^z} be the set of $i \in I$ such that Q_t^i is consistent with $[\chi_t(A) = z]$. Since we are assuming that $[\chi_t(A) = z]$ and E_t consist entirely of t -admissible information, they must be entailed by Q_t^i for each $i \in I_{A^z}$.

We begin by verifying that if $i \in I_{A^z}$, then $c_G(A|Q_t^i) = z$. Let $i \in I_{A^z}$, and suppose Q_t^i is true at world w . Since Q_t^i entails $[\chi_t(A) = z]$, $[\chi_t(A) = z]$ is also true at w . But Chance as Evidential Support tells us that for $[\chi_t(A) = z]$ to be true at w is for $sup(A|Q_t^i) = z$ to be true at w . It follows that Q_t^i must entail $sup(A|Q_t^i) = z$, and, therefore that $c_G(A|Q_t^i) = c_G(A|Q_t^i \cdot [sup(A|Q_t^i) = z])$. One can then use the Evidential Principle to derive $c_G(A|Q_t^i) = z$.

One can then derive the Principal Principle (relativized to G), as follows:

$$\begin{aligned}
c_G(A|E_t[\chi_t(A) = z]) &= \sum_{i \in I_{A^z}}^{\infty} c_G(A|E_t[\chi_t(A) = z]Q_t^i) \cdot c_G(Q_t^i|E_t[\chi_t(A) = z]) && \text{count. additivity} \\
&= \sum_{i \in I_{A^z}}^{\infty} c_G(A|E_tQ_t^i) \cdot c_G(Q_t^i|E_t[\chi_t(A) = z]) && Q^i \text{ entails } [\chi_t(A) = z] \\
&= \sum_{i \in I_{A^z}}^{\infty} c_G(A|Q_t^i) \cdot c_G(Q_t^i|E_t[\chi_t(A) = z]) && Q_t^i \text{ entails } E_t \\
&= \sum_{i \in I_{A^z}}^{\infty} z \cdot c_G(Q_t^i|E_t[\chi_t(A) = z]) && c_G(A|Q_t^i) = z \\
&= z && \text{def. of } I_{A^z}
\end{aligned}$$

References

- Arntzenius, F., Elga, A., & Hawthorne, J. (2004), ‘Bayesianism, infinite decisions, and binding’, *Mind* **113**, 251–283.
- Braddon-Mitchel, D. & Nola, R., eds (2009), *Conceptual Analysis and Philosophical Naturalism*, MIT Press, Cambridge, MA.
- Carnap, R. (1962), *Logical Foundations of Probability*, University of Chicago Press, Chicago.
- Carnap, R. (1971), ‘Inductive logic and rational decisions’. In Carnap & Jeffrey (1971), pp. 5–31.
- Carnap, R. & Jeffrey, R. C., eds (1971), *Studies in inductive logic and probability*, Vol. 1, University of California Press, Berkeley.
- Christensen, D. (2010), ‘Rational reflection’, *Philosophical Perspectives* **24**, 121–140.
- Dummett, M. (1991), *The Logical Basis of Metaphysics*, Harvard, Cambridge, MA.
- Easwaran, K. (2013), ‘Why countable additivity?’, *Thought* **1**, 53–61.
- Easwaran, K. (forthcoming), ‘Conditional probability’. In Hájek & Hitchcock (forthcoming).
- Eriksson, L. & Hájek, A. (2007), ‘What are degrees of belief?’, *Studia Logica: An International Journal for Symbolic Logic* **86**, 183–213.
- Field, H. (2008), *Saving Truth from Paradox*, Oxford University press, Oxford.
- Hájek, A. & Hitchcock, C., eds (forthcoming), *The Oxford Handbook of Probability and Philosophy*, Oxford University Press, Oxford.
- Hall, N. (1994), ‘Correcting the guide to objective chance’, *Mind* **103**, 505–518.

- Hempel, C. G. (1945*a*), ‘Studies in the logic of confirmation (i.)’, *Mind* **54**(213), 1–26.
- Hempel, C. G. (1945*b*), ‘Studies in the logic of confirmation (ii.)’, *Mind* **54**(214), 97–121.
- Hill, B. M. & Lane, D. (1985), ‘Conglomerability and countable additivity’, *Sankhyā: The Indian Journal of Statistics, Series A*, **47**, 366–379.
- Holton, R. (forthcoming), ‘Intention as a model for belief’. In Vargas & Yaffe (forthcoming).
- Jeffrey, R. C., ed. (1980), *Studies in Inductive Logic and Probability*, Vol. 2 of *Proceedings of the Symposia on Pure Mathematics*, University of California, Berkeley.
- Kadane, J. B., Schervish, M. J. & Seidenfeld, T. (1999), *Rethinking the Foundations of Statistics*, Cambridge Studies in Probability, Induction and Decision Theory, Cambridge University Press.
- Kolmogorov, A. N. (1960), *Foundations of the Theory of Probability*, 2 edn, Chelsea Pub Co.
- Lewis, D. (1980), ‘A subjectivist’s guide to objective chance’. In Jeffrey (1980). Reprinted with added postscripts in Lewis (1986*b*), pp. 83–132.
- Lewis, D. (1983), ‘New work for a theory of universals’, *Australasian Journal of Philosophy* **61**, 343–377. Reprinted in Lewis (1999).
- Lewis, D. (1986*a*), *On the Plurality of Worlds*, Blackwell, Oxford and New York.
- Lewis, D. (1986*b*), *Philosophical Papers, Volume II*, Oxford.
- Lewis, D. (1994), ‘Humean supervenience debugged’, *Mind* **103**, 473–490. Reprinted in Lewis (1999).
- Lewis, D. (1999), *Papers in Metaphysics and Epistemology*, Cambridge University Press, Cambridge.

- Lewis, D. (2009), ‘Ramseyan humility’. In Braddon-Mitchel & Nola (2009).
- McGee, V. (1999), ‘An airtight dutch book’, *Analysis* **59**, 257–8211.
- McManus, D. (2004), *Wittgenstein and Scepticism*, Routledge, London and New York.
- Priest, G. (2006), *In Contradiction: A study of the Transconsistent*, second edn, Oxford University Press, Oxford.
- Rayo, A. (2013), *The Construction of Logical Space*, Oxford University Press, Oxford.
- Schervish, J. B., Kadane, M. J. & Seidenfeld, T. (1999), ‘Statistical implications of finitely additive probability’, pp. 211–231. In Kadane et al. (1999).
- Schervish, M. J., Seidenfeld, T. & Kadane, J. B. (1984), ‘The extent of non-conglomerability of finitely additive probabilities’, *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **66**, 205–226.
- Seidenfeld, T. & Schervish, M. (1983), ‘A conflict between finite additivity and avoiding dutch book’, *Philosophy of Science* **50**, 398–412.
- Skyrms, B. (1986), *Choice and Chance*, third edn, Wadsworth, Belmont, CA.
- Stalnaker, R. C. (typescript), ‘Counterfactuals and humean reduction’.
- Titelbaum, M. G. (2010), ‘Not enough there there: Evidence, reasons, and language independence’, *Philosophical Perspectives* **24**, 477–528.
- Titelbaum, M. G. (2011), ‘Symmetry and evidential support’, *Symmetry* **3**, 680–698.
- Vargas, M. & Yaffe, G., eds (forthcoming), *Rational and Social Agency: Essays on the Philosophy of Michael Bratman*, Oxford University Press, Oxford.
- Vitali, G. (1905), ‘Sul problema della misura dei gruppi di punti di una retta’, *Bologna, Italy* .

- White, R. (2005), 'Epistemic permissiveness', *Philosophical Perspectives* **19**.
- White, R. (2010), 'Evidential symmetry and mushy credence', *Oxford Studies in Epistemology* **3**. Tamar Szabo Gendler and John Hawthorne eds.
- Williamson, J. (1999), 'Countable additivity and subjective probability', *British Journal for the Philosophy of Science* **50**, 401–416.
- Wittgenstein, L. (1969), *On Certainty*, Blackwell, Oxford. Edited by Anscombe, G.E.M. and von Wright, G.H. Translated by Anscombe, G.E.M. and Paul, D.
- Wright, C. (2001), 'On being in a quandary: Relativism, vagueness, logical revisionism', *Mind* **110**, 45–98.
- Wright, C. (2004), 'Wittgensteinian certainties'. In McManus (2004).