Impermissive Bayesianism

Christopher J. G. Meacham

October 13, 2013

Abstract

This paper examines the debate between permissive and impermissive forms of Bayesianism. It briefly discusses some considerations that might be offered by both sides of the debate, and then replies to some new arguments in favor of impermissivism offered by Roger White. First, it argues that White's (2010) defense of Indifference Principles is unsuccessful. Second, it contends that White's (2005) arguments against permissive views do not succeed.

1 Introduction

At the heart of Bayesianism is a pair of normative constraints: an agent's degrees of belief should satisfy the probability axioms, and an agent should update her beliefs via conditionalization. But there's a debate among Bayesians concerning whether there are any further normative constraints, and if so, how strong these additional constraints are. *Impermissive Bayesians* take there to be additional constraints strong enough to uniquely fix what the beliefs of any agent with a given body of evidence should be. *Permissive Bayesians* take there to be additional constraints that are weaker than this, or take there to be no additional constraints at all.

Several Bayesians, such as Christensen (2007) and White (2005), have recently come out in favor of impermissive Bayesianism. In this paper, I'll review some considerations that might be taken to favor one side or the other, and assess some of the new arguments for impermissivism that have appeared in the recent literature. In particular, I'll focus on the arguments presented in White (2005) and White (2010).

I'll proceed as follows. In the next section I'll sketch some background. In the third section I'll briefly look at four considerations that might be offered for or against each side of the debate: cases invoking permissive intuitions, cases invoking impermissive intuitions, principles that constrain rational belief, and worries regarding Indifference Principles and their relatives. I'll suggest that these considerations are largely inconclusive. I'll then turn to examine some recent attempts to bolster impermissivism offered by Roger White. In the fourth section I'll consider White's (2010) defense of Indifference Principles, and I'll argue that it does not succeed. In the fifth section I'll assess four of White's (2005) arguments against permissivism, and I'll contend that these arguments are unsuccessful. In the sixth section I'll conclude with some brief remarks.

2 Background

Let us characterize an agent's doxastic state with a function that assigns real numbers to propositions. These values, called *credences* or *degrees of belief*, represent the agent's confidence in a proposition, where a credence of 1 indicates that she is virtually certain the proposition is true, and a credence of 0 indicates that she is virtually certain the proposition is false.[1]

For the purposes of this paper, I will take an account of epistemic norms to be *Bayesian* if it entails the following two normative constraints:[2]

Probabilism: An agent's credences cr should satisfy the probability axioms.

Conditionalization: If an agent with credences cr receives some new evidence E, then her new credences cr_E should be:

$$cr_E(A) = cr(A|E), \text{ if defined.}$$

There are a number of ways to carve up the space of Bayesian positions. Some of these ways are more fine-grained than others. In preparation for the discussion to come, it will be useful to go through a few of these different ways of dividing things up.

First, we can broadly divide Bayesian accounts into two camps, depending on whether they entail the following condition:[3]

Evidential Uniqueness: For any evidence E, there is a unique doxastic state that any agent with total evidence E should have.

Impermissive Bayesian accounts entail Evidential Uniqueness. *Permissive Bayesian* accounts do not.[4]

Consider the initial credences, or "priors", that an agent might have prior to receiving any evidence.[5] We can divide Bayesian accounts in a more fine-grained way by classifying them according to which priors they take to be permissible. To set things up cleanly, let us restrict

[1] Only *virtually* certain because in each case there could be measure 0 exceptions.

[2] Although this tracks one use of the term "Bayesian", this term has been used in a number of different ways. Some have taken Bayesians to only be committed to Probabilism or only committed to Conditionalization. Likewise, Probabilism and Conditionalization themselves have been understood in slightly different ways. For example, some have restricted Conditionalization to agents with probabilistic credences, and Probabilism to agents whose doxastic states are fine-tuned enough to admit of precise degrees.

[3] I take this condition from White (2005) and Feldman (2007), which they call "Uniqueness". I've called it "Evidential Uniqueness" in order to emphasize the key feature of the thesis—that one's doxastic state should be uniquely fixed by all and only one's evidence.

[4] The distinction between permissive and impermissive Bayesianism mirrors the distinction between subjective and objective Bayesianism. I've avoided the subjective/objective terminology because these terms have been used in a number of different ways. And while pretty much every way of making this distinction would classify impermissive Bayesianism as a form of objective Bayesianism, there's little consensus beyond that with respect to where the line between subjective and objective Bayesianism should be drawn.

[5] I'm assuming here a picture of evidence according to which agents start with no evidence. But nothing hangs on this—those inclined to favor a different picture of evidence can just understand my talk of "evidence" to mean "non-initial evidence".

our attention from now on to cases in which agents have such initial credence functions, and only get E as evidence if they already have a non-zero credence in E.[6]

If an agent satisfies conditionalization, her previous credences cr and her new evidence E completely determine her new credences cr_E. And each previous credence function is fixed by the credence function and evidence before that, all the way back to her initial credence function. So if an agent satisfies conditionalization, her priors and the evidence she's gotten will completely fix her new credences.

Now consider the range of permissible priors.[7] One can see the debate between permissive and impermissive Bayesians as a debate about how broad this range of priors is. The extreme permissive Bayesians maintain that all (probabilistically coherent) priors functions are permissible. The impermissive Bayesians maintain that there's only one permissible priors function. And there's a broad range of views in-between these two extremes, which vary with respect to how broad they take the range of permissible priors to be.

So here is one way to divide up the space of Bayesian accounts:[8]

Extreme Permissive Bayesian: Bayesian accounts which maintain that any probabilistic priors function is rationally permissible.

Moderate Permissive Bayesian: Bayesian accounts which maintain that more than one probabilistic priors function is rationally permissible, but not all of them are.

Impermissive Bayesian: Bayesian accounts which maintain that only one probabilistic priors function is rationally permissible.

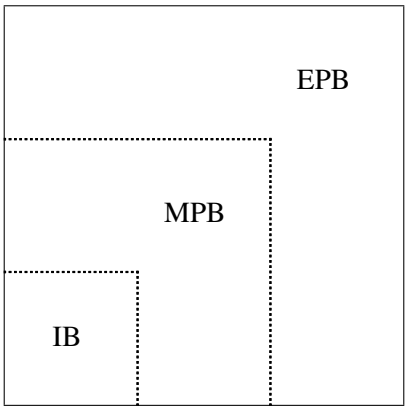

Bayesian Views

[6]So we can ignore agents with infinite pasts or agents who exist during open intervals of time, and ignore cases where $cr(A|E)$ is undefined.

[7]I.e., the range of priors functions held by possible rational agents.

[8]The diagram below should be understood as representing the range of Bayesian views, with different points in the space corresponding to different Bayesian accounts. (Impermissive Bayesianism is not a single point because there are many different versions of impermissive Bayesianism, which differ with respect to what they take the one permissible priors function to be.)

Here is another way to divide up Bayesian accounts that will be useful in what follows. Note that, given Bayesianism, Evidential Uniqueness is true *iff* both of the following claims are true:

Agent Uniqueness: For any possible agent, there is only one permissible priors function.

Permission Parity: The same priors functions are permissible for all possible agents.

To see that this equivalence holds, first suppose that Bayesianism and Evidential Uniqueness are true. Consider an agent who has just been created *tabula rasa*, and so has no evidence. Evidential Uniqueness requires there to be only one permissible doxastic state for such an agent. So Agent Uniqueness is true. Furthermore, Evidential Uniqueness requires this state to be the same for all possible *tabula rasa* agents. (If this wasn't the case, then the evidence alone wouldn't suffice to fix what an agent ought to believe. We would also need some further facts, such as facts about who the agent is, or what her world is like, in order to fix what she ought to believe. And if this were the case, then Evidential Uniqueness would be false.) So Permission Parity is true.

Going the other way, suppose that Bayesianism, Agent Uniqueness and Permission Parity are true. Agent Uniqueness and Permission Parity entail that there is only one permissible priors function for all possible agents. And given total evidence E, an agent's priors and conditionalization will pick out a single permissible doxastic state. So for any total evidence E, there is a unique permissible doxastic state that an agent with that evidence should have. So Evidential Uniqueness is true.[9]

So given Bayesianism, Evidential Uniqueness is true *iff* Agent Uniqueness and Permission Parity are true. Since permissive Bayesians reject Evidential Uniqueness, they must reject Agent Uniqueness or Permission Parity. This gives us another way to divide up the space of Bayesian accounts:

Impermissive Bayesian: Bayesian accounts which accept both Agent Uniqueness and Permission Parity, holding that there is only one permissible priors function for each agent, and that which priors function is permissible is the same for all possible agents.

Permissive$_1$ Bayesian: Bayesian accounts which reject Agent Uniqueness but accept Permission Parity, holding that there are multiple permissible priors functions for each agent, and that which priors functions are permissible is the same for all possible agents.

Permissive$_2$ Bayesian: Bayesian accounts which reject both Agent Uniqueness and Permission Parity, holding that there can be multiple permissible priors functions for an agent, and that which priors functions are permissible varies among possible agents.

Permissive$_3$ Bayesian: Bayesian accounts which accept Agent Uniqueness but reject Permission Parity, holding that there is only one permissible priors function for each agent, and that which priors function is permissible varies among possible agents.

[9]Note that the argument won't go through if we don't put aside cases with 0-credence evidence and agents without initial credences. If we don't put aside such cases, the entailment will only go one way: given Bayesianism, Evidential Uniqueness will entail Agent Uniqueness and Permission Parity, but not vice versa.

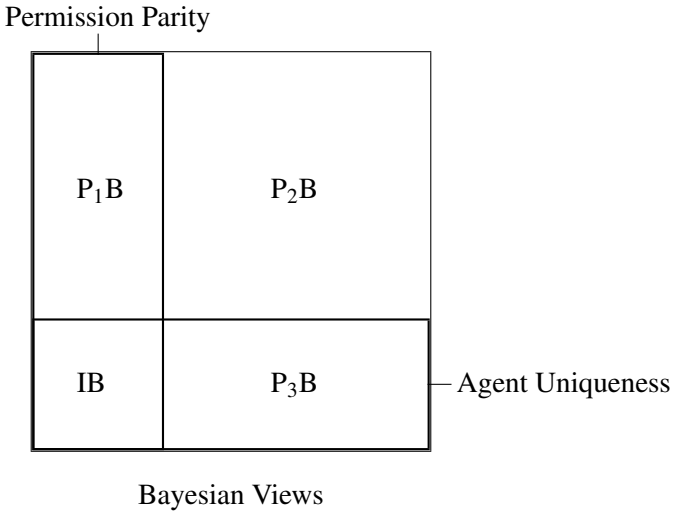

Together with the distinctions we drew earlier, we can lay out the Bayesian landscape as follows:

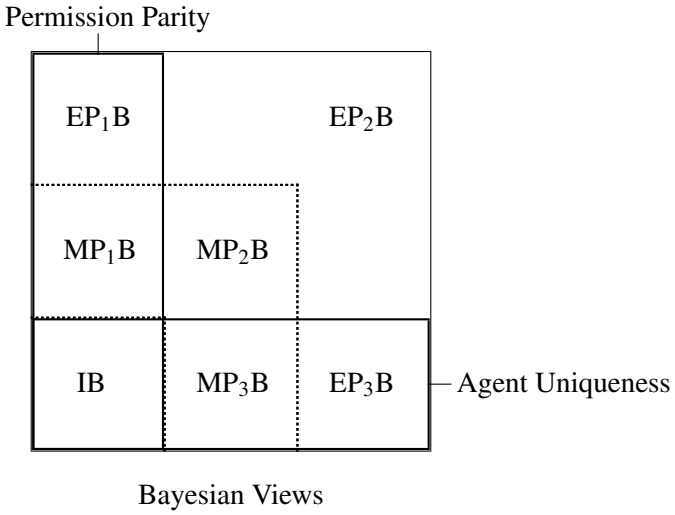

3 Four Considerations

3.1 Permissive Intuitions

One source of potential objections to impermissive Bayesianism are cases which invoke permissive intuitions regarding what an agent should believe. These cases might include jury verdicts (should you believe the defendant is guilty?), moral beliefs (should you be a Utilitarian?), and theological beliefs (should you believe in God?). But some of the most compelling cases involve the initial beliefs of agents. In these cases, the objection might go something like this:

> Consider an agent who has just been created *tabula rasa*. Consider her initial belief state. At first glance, it's implausible to think that there's a single doxastic state that she should be rationally required to have. It's hard to see how rationality could require her to have a particular credence in (say) the proposition that there are several thousand chickens nearby. After all, what would this unique rational credence be? 0.7? 0.1? Given no evidence and no background beliefs to appeal to, how could one think that her credence is rationally required to take any one of these values?

These cases have considerable intuitive force. But it's worth getting clear on what intuitions are in play here, and what positions they support.

Here are two permissive intuitions that might be invoked by these cases:

Permissive Intuition 1: There are evidential situations in which two different agents can rationally adopt different beliefs.

Permissive Intuition 2: There are evidential situations in which a particular agent can rationally adopt a range of different beliefs.

The first intuition is that there's flexibility with respect to what different agents in a particular evidential situation are permitted to believe. While the permissive Bayesian can accommodate this intuition, the impermissive Bayesian cannot. So insofar as we have this intuition, we have a reason to favor permissive over impermissive Bayesianism.

The second intuition is that there's flexibility with respect to what a particular agent is permitted to believe. And for the most part neither permissive nor impermissive Bayesianism can accommodate this intuition. More precisely, we can divide this intuition into two parts:

Permissive Intuition 2a: There are *tabula rasa* cases in which a particular agent can rationally adopt a range of different beliefs.

Permissive Intuition 2b: There are non-*tabula rasa* cases in which a particular agent can rationally adopt a range of different beliefs.

The first part of this intuition is compatible with versions of Bayesianism that reject Agent Uniqueness, since these views allow rational agents to begin with a range of initial belief states. But the second part of this intuition is incompatible with any kind of Bayesianism. All Bayesians hold that an agent's prior beliefs and her new evidence suffice to fix what her new beliefs should be. So in non-*tabula rasa* cases, what a particular agent ought to believe is fixed.

So when evaluating the bearing of permissive intuitions on the Bayesian permissive/impermissive debate, we should keep in mind that cases like those described above can be used to invoke several different kinds of permissive intuitions. Only some of these intuitions, like 1 and 2a, favor some form of permissive Bayesianism over impermissive Bayesianism. And none of these intuitions give us a reason to favor extreme permissive Bayesianism over moderate permissive Bayesianism.

3.2 Impermissive Intuitions

One source of potential objections to permissive Bayesianism are cases which invoke impermissive intuitions regarding what an agent should believe. An example of such a case is the grue/green case discussed by Goodman (1954):

> Let something be *grue iff* it is green and has been observed before 2020, or if it is blue and has not been observed before 2020. Consider the following three propositions:
>
> E**:** Every emerald observed before 2020 is green.
>
> H_1**:** All emeralds are green.
>
> H_2**:** All emeralds are grue.
>
> Suppose that we have some non-extremal credence in each of these propositions, and that we know nothing about emeralds or what color they might be. And suppose we now get E as evidence. Both H_1 and H_2 entail E. But surely upon receiving this evidence we ought to believe H_1 is true, not H_2.

Cases of this kind are intuitively compelling. But it's worth paying attention to what intuitions are being appealed to, and what views they tell in favor of.

Consider three intuitions one might have about the case above:

Impermissive Intuition 1: After receiving evidence E, an agent should increase her credence in H_1, but not H_2.

Impermissive Intuition 2: After receiving evidence E, an agent should have a higher credence in H_1 than H_2.

Impermissive Intuition 3: After receiving evidence E, an agent like us (i.e., an agent who is cognitively similar to us, with priors similar to ours) should have a higher credence in H_1 than H_2.

The first intuition is that only H_1 should be confirmed by E. This is the result that Goodman (1954) himself wanted. But this intuition does not favor impermissive over permissive Bayesianism, because neither view can deliver this result. Bayesians of any stripe maintain that if H entails E, then an agent's credence in H should go up upon receiving E as evidence.[10] Since both H_1 and H_2 entail E, a Bayesian will maintain that an agent's credence in both hypotheses should go up.

The second intuition is that any rational agent should end up thinking that H_1 is more likely than H_2. Impermissive Bayesians can impose rationality constraints which require this, as can moderate permissive Bayesians. Extreme permissive Bayesians, on the other hand, must allow some agents to rationally believe otherwise. So insofar as we have this intuition, we have a reason to favor impermissive or moderate permissive Bayesianism over extreme permissive Bayesianism.

The third intuition is that any rational agent like us should end up thinking that H_1 is more likely than H_2. This intuition does not favor impermissive over permissive Bayesianism,

[10] Assuming, as in the case above, that the agent initially has a non-extremal credence in both H and E.

because both views can yield this result. Given the cognitive similarity between different humans, it's plausible that typical humans will have similar priors. If typical humans have similar priors, then Bayesians will maintain that typical humans should adopt similar beliefs given evidence E. So any Bayesian can maintain that agents like us should come to have a higher credence in H_1 than H_2 in the case described above.

So when assessing the influence of impermissive intuitions on the Bayesian permissive/impermissive debate, we should remember that these kinds of cases can be used to invoke several different kinds of impermissive intuitions. Many of these intuitions, like 1 and 3, are orthogonal to the permissive/impermissive debate. And none of these intuitions favor impermissive Bayesianism over moderate permissive Bayesianism.

3.3 Plausible Principles

Another potential source of objections to permissive Bayesianism appeals to principles that constrain rational belief.[11]

The most widely accepted of these principles are Chance-Credence Principles.[12] Chance-Credence Principles require that a rational agent's beliefs line up with what she thinks the chances are in certain ways. For example, if the agent's evidence consists only of the fact that A has a 50% chance of becoming true, then she should have a credence of 0.5 in A. Chance-Credence Principles constrain what rational initial credence functions are permissible. Thus these principles are incompatible with extreme permissive Bayesianism. And since it is plausible that some such principle obtains, this gives us a strong reason to reject extreme permissive Bayesianism. That said, Chance-Credence Principles allow for a broad range of permissible priors. So although these principles tell against extreme permissive Bayesianism, they don't tell against moderate permissive Bayesianism.

A number of other principles constraining rational belief have been proposed.[13] But the Chance-Credence Principle alone already rules out extreme permissive Bayesianism. And adding these other principles to the Chance-Credence Principle still doesn't yield a constraint strong enough to rule out moderate permissive Bayesianism. So given that some Chance-Credence Principle is correct, these other constraints add little to the permissive/impermissive debate.

There is one exception. Let a *Strong Indifference Principle* be an Indifference Principle that picks out a single permissible priors function.[14] If one of these strong Indifference Principles is correct, then all forms of permissive Bayesianism are false. So a lot hangs on the tenability of these Indifference Principles. (There are also weaker Indifference Principles, which can be

[11] Since these objections rely on the plausibility of certain principled constraints on rational belief over and above those imposed by Bayesianism, these objections can be seen as a special case of the impermissive intuitions objections.

[12] For classic discussions, see Lewis (1986), Hall (1994) and Lewis (1994).

[13] In addition to Chance-Credence Principles and Indifference Principles (described below), these proposals include Reflection Principles (see van Fraassen (1984)), Expert Principles (see Gaifman (1988)) and Regularity Principles (see Howson (2000)).

[14] These principles have been given a number of different names, including "The Principle of Indifference", "The Principle of Insufficient Reason", "The Maximum Entropy Principle", and "Jeffrey's Rule". (See Howson and Urbach (2005), and the references therein.)

grouped with the other constraints on rational belief discussed above. As these weak Indifference Principles add little to the impermissive/permissive debate, I'll restrict my attention to strong Indifference Principles in what follows.) We turn to this topic next.

3.4 Indifference Principles

The viability of Indifference Principles bears on the permissive/impermissive debate in two ways. First, given Bayesianism, Indifference Principles entail that some form of impermissive Bayesianism must be true, since they pick out a single rational initial credence function. Second, it's generally been thought that one needs something like an Indifference Principle in order to obtain rationality constraints that are as strong as impermissive Bayesianism requires. Thus most proponents of impermissive Bayesianism have championed some form of Indifference Principle.

Indifference Principles take the following form:

Indifference Principle: If an agent is "indifferent" with respect to some set S of mutually exclusive propositions, then her credence in S should be "appropriately distributed" among the members of S.

The terms "indifferent" and "appropriately distributed" serve as placeholders for more substantive claims. By filling in these placeholders in different ways, we obtain different Indifference Principles.

Why adopt an Indifference Principle? The usual reasons stem from a combination of abstract intuitions about how one ought to be "appropriately unbiased", and particular intuitions about cases.[15]

The abstract intuitions stem from the thought that our initial beliefs should not unfairly favor one empirical hypothesis over another. The idea is that an adequate account of how to respond to evidence should be neutral and "let the data speak for itself". On a Bayesian account, this requires having priors that are appropriately unprejudiced and even-handed about the significance of the evidence.

The particular intuitions appeal to cases in which indifference-style prescriptions are plausible. For example, given that you know only that one of the three doors in front of you has a prize behind it, what should your credence be that the prize lies behind the first door? Many have the intuition that your credence should be 1/3. Likewise, if you know only that it is Monday or Tuesday, many have the intuition that your credence that it is Monday should be 1/2.[16] And there are a large number of cases like this. Since these "indifferent" verdicts are intuitively plausible ones, this is taken to support the idea that Indifference Principles encode rationality constraints.

[15] Early defenders of Indifference Principles also appealed to empirical evidence to support these principles (see Jaynes (1983)). For example, it was noted that the degrees of belief suggested by some Indifference Principles matched the real-world frequencies of various thermodynamic phenomena, and it was suggested that this gave us a reason to believe these Indifference Principles were true. But this is now widely recognized by both proponents and opponents of Indifference Principles to be a mistake (see North (2010), White (2010)).

[16] Both of these cases come from White (2010).

What reasons are there to worry about Indifference Principles? A number of specific worries arise for different versions of these principles.[17] But there are also some general worries that arise for all Indifference Principles. I'll focus on these general worries here.

Consider a case offered by van Fraassen (1989):

The Cube Factory: You know a particular factory produces cubes whose height in centimeters lies in the (0,2] interval.[18] Furthermore, you know that a particular cube produced by this factory has been selected. You know nothing more about the factory or how the cube was selected.

Given this case, what should your credence be that this cube is (0,1] cm versus (1,2] cm high? When presented with this question, we have the intuition that you should be indifferent between the two possibilities, and assign each a credence of 1/2.

But there are other ways of describing the case that invoke conflicting intuitions. The cubes that this factory produces will have a face area that lies in the (0,4] cm^2 interval. Given this, what should one's credence be that the faces of this cube are (0,2] cm^2 in area versus (2,4] cm^2 in area? When presented with this question, we again have the intuition that you should be indifferent between the two possibilities, and assign each a credence of 1/2. But these two verdicts are inconsistent: if one's credence is evenly split between the cube being (0,1] and (1,2] cm high, then one's credence should be evenly split between the cube's face area being (0,1] and (1,4] cm^2, not (0,2] and (2,4] cm^2.

It is sometimes suggested that this tension between different plausible prescriptions only arises for certain special cases, like the cube factory case. And there are plenty of other cases, such as the three door and Monday/Tuesday cases described above, for which these problems do not arise.[19] I think this is a mistake—similar problems arise for all of these cases. With this in mind, I'll sketch a number of other ways in which one might be indifferent in the cube factory case. This will give us a better feel for the variety of ways in which we might plausibly be indifferent, and make it easier to see how the same issues arise in other cases where these problems are less apparent.

We've considered being indifferent with respect to height and face area, but there are a number of other quantities one might be indifferent with respect to. For example, one might be indifferent with respect to the inverse height of the cube. Or one might be indifferent with respect to some gruesome quantity of cubes, g, where a cube's g is equal to the distance between opposing corners if it's height lies in the $(0, \sqrt[4]{3}]$ interval, and equal to its volume otherwise. These other quantities will yield more conflicting prescriptions.

[17] For a survey of the particular problems that confront various attempts to flesh out a Principle of Indifference, see Howson and Urbach (2005) and Weisberg (2011).

[18] $(x, y]$ is the interval between x and y that includes y but not x.

[19] For example, White (2010) takes the moral of the cube factory case to be that there will be some tricky cases, like the cube factory case, in which we won't know how to apply the Indifference Principle. But, White maintains, this doesn't make the Indifference Principle useless, as there are plenty of other cases, such as the three door and Monday/Tuesday cases, in which it is obvious how one should be indifferent (see White (2010), p.167-169). I want to suggest, in these other cases, that it's only obvious how to be indifferent in the same sense that it's obvious, upon first being presented with the cube factory case, that you should be indifferent with respect to length. While it may at first seem obvious how we should be indifferent, further reflection should make our confidence in these verdicts evaporate.

Suppose we settle on a given quantity, say height. There are various units one might use to measure height. One might measure height in meters. Or one might measure height in holdons, where the height of an object in holdons is equal to the $\log_{10}$ of its height in meters. These choices will yield incompatible prescriptions. And there are infinitely many other scales that are non-linearly related to one another to choose from.

Here is another way in which to be indifferent. Consider the possible arrangements of particles in the universe. To simplify a bit, let us suppose that we've fixed the spatiotemporal features of the world and the number of particles.[20] (If anything, this simplification makes things easier for the proponent of indifference, since all of these worries re-arise when we consider how to be indifferent over the different spatiotemporal features the world could have and the different numbers of particles there could be.) Given this, we can represent the space of possible configurations of these particles using some high dimensional space in which each point corresponds to a possible configuration, and each degree of freedom corresponds to a different dimension.[21] One way to be indifferent is with respect to some measure μ over this high dimensional space.

Now, some of these configurations will be compatible with the existence of a cube factory which produces cubes (0,2] cm in height and with one of these cubes having been selected and used to query an agent like yourself. These possibilities will select some region R of our high dimensional space. And some of the possibilities in R will be ones in which the selected cube is (0,1] cm in height; call this region S. Given the indifference assignment suggested above, one's credence that the cube is (0,1] cm in height should be equal to the proportion of R taken up by S according to μ.

What values this procedure recommends will depend on what other choices we make, such as which quantities we choose to represent the various degrees of freedom and which natural measure over these quantities we choose (should the measure be linear in meters or holdons?). Different choices will yield different prescriptions. And these prescriptions will generally conflict with both each other and the other prescriptions described above.

Similar worries arise for every case of indifference. For example, consider the Monday/Tuesday case described above, where we're initially inclined to divide our credence equally between the two days. While we get this prescription if we're indifferent in a manner that's uniform with respect to time-in-seconds, uniformity with respect to different choices of quantities or units will yield different results.[22] Likewise, in the three door case described above,

[20]Note that some assumptions of this kind are already assumed in the initial descriptions of the case. For example, the height and face area descriptions of the case together require space to have at least three dimensions, to be large enough to hold at least a 2 cm cube and a factory for making such cubes, to have features which yield the straightforward relationship between length and face area we're accustomed to, and so on.

[21]I assume here that the number of particles is finite, so that the number of degrees of freedom in the system is finite. I also assume here that the spatiotemporal extension of the world is finite. (Again, these simplifications make things easier for the proponent of indifference, since it allows them to avoid various infinity worries.)

[22]For example, one might be indifferent in a manner that's uniform with respect to something like inverse time, assessed with respect to the beginning of the universe. (E.g., let t = the number of seconds since the beginning of the universe, let $\tau = \frac{1}{t+1}$, and require rational priors to be uniform with respect to τ. (This normalizes because we're assuming the spatiotemporal extension of the world is finite; see footnote 21.)) Since Tuesday is later than Monday, the difference in inverse time between the beginning of Tuesday and the end of Tuesday will be smaller than the

being indifferent over initial conditions with respect to some choice of quantities and measure will generally not yield the result that one's credence that the prize is behind the first door should be 1/3.[23]

These kinds of problems point to three general worries for Indifference Principles. The first worry is that proponents of Indifference Principles face a trilemma, each horn of which appears problematic. The second worry is that there's little to be gained by adopting an Indifference Principle. The third worry is that the indifference intuitions that motivate these principles are untrustworthy. Let's look at each of these worries in turn.

The first worry is that proponents of Indifference Principles face a trilemma, and each horn of this trilemma is unappealing. Let's begin by looking at how this worry arises in the cube factory case. In this case, we have the intuition that one ought to be indifferent with respect to height, but also the intuition that one ought to be indifferent with respect to face area, and so on. One option is to require indifference in all of these respects. But since these different ways of being indifferent conflict, this leads to inconsistent prescriptions. A second option is to not require rational agents to be indifferent in any particular respect, but to just be indifferent in some respect or other, whether it be height, face area, or something else. But since there are any number of ways in which one can be indifferent, this leads to trivial prescriptions. A third option is to pick one of these ways of being indifferent—height, say—and require rational agents to be indifferent in this respect. But since there doesn't seem to be any good reason to choose one of these ways over the others, this leads to arbitrary prescriptions.

More generally, every Indifference Principle depends on how we carve up or represent the space of possibilities.[24] Proponents of these principles can handle this dependency in one of three ways:

Option 1: Require rational agents to be indifferent with respect to *every* carving.

Option 2: Allow rational agents to be indifferent with respect to *any* carving.

Option 3: Require rational agents to be indifferent with respect to one particular carving.

All three options appear problematic. The first option appears to yield inconsistent principles. The second option appears to yield trivial principles. And the third option appears to yield arbitrary principles.[25]

---

corresponding interval for Monday, and Monday will be assigned a higher value.

[23]For example, consider the spatial dimension along which the doors are arranged. Consider a configuration space in which n of the degrees of freedom correspond to the positions of each of the n particles along this spatial dimension, and where one is indifferent in a manner that's uniform over something like inverse distance along this dimension, assessed with respect to some point to the left of the three doors. (E.g., let x = the number of meters away from the point along this spatial dimension in the relevant direction, let $\xi = \frac{1}{|x|+1}$, and require rational priors to be uniform with respect to ξ.) Then the credence assigned to the prize being behind each door, as we go from left to right, will decrease.

[24]The ways in which this dependence arises depends on how one spells out the Indifference Principle. For a discussion of some of these details, see Howson and Urbach (2005) and Weisberg (2011).

[25]Note that finding a carving that's not arbitrary in some respect is not what is required to avoid the arbitrariness horn of the trilemma. What is required is to find a carving that makes the resulting principle *epistemically* non-arbitrary. One might make a case that some particular carving is more natural than the others in some salient respect—perhaps it lines up with the perfectly natural properties, for example (see Schaffer (2007), Sider (2011)). But the existence of such a carving doesn't by itself give us any reason to think that an Indifference Principle should employ it. After all,

The second worry is that there's little intuitive payoff to be gained by adopting an Indifference Principle. Grant for the sake of argument that indifference intuitions provide good evidence regarding what we ought to believe. Proponents of indifference generally motivate Indifference Principles by presenting a number of cases in which we have indifference-style intuitions, such as the three door and Monday/Tuesday cases described above. But these cases do little to support the adoption of an Indifference Principle until we've tied these intuitive cases of indifference to a single, consistent principle. Until then there isn't any reason to think that there's a single rationality constraint which these intuitions are supporting, as opposed to (say) a number of mutually inconsistent principles that yield one or two of these intuitive verdicts, and conflict with all of the rest.

These worries are borne out by our conflicting indifference judgments. We have a large number of mutually inconsistent indifference intuitions. And no consistent principle can capture more than a sliver of them. So the fact that we have a number of strong indifference intuitions doesn't give us much reason to adopt an Indifference Principle. Because almost all of these indifference intuitions are going to have to be rejected, regardless of whether we adopt such a principle.

The third worry is that the inconsistency of our indifference intuitions gives us reason to doubt that they're trustworthy. After all, the problems in the cube-factory case don't arise from poor formulations of Indifference Principles. The problems arise because our indifference intuitions themselves are inconsistent. And this gives us good reason to believe these intuitions are unreliable.[26]

In light of these kinds of worries, one can see why Indifference Principles have fallen out of favor. But White (2010) has recently argued that this poor opinion of Indifference Principles is unjustified, and has defended a version of the Indifference Principle.[27] We turn now to White's defense.

4 White's Case for Indifference

White (2010) proposes the following Indifference Principle:

White's Indifference Principle: If A and B are evidentially symmetric for an agent, then her credence in A and B should be equal.

---

the existence of carvings that are non-arbitrary (in this sense) is compatible with there being no constraints on rational belief at all.

[26] Of course, a proponent of Indifference Principles who did not appeal to indifference intuitions to motivate the adoption of these principles would not be subject to the second and third worries given above. Instead, they would face the challenge of finding some other compelling reason for adopting an Indifference Principle.

[27] As we will see, it is unclear whether White intends for his principle to be a strong Indifference Principle, and thus unclear whether he takes it to be a principle one could use to support Evidential Uniqueness. Given the understanding of White I'll suggest, White is neutral with respect to whether his principle is a strong Indifference Principle or not (see section 4.1). But since, as I understand him, he is amenable to it being a strong Indifference Principle, we will need to examine it in order to see whether he has found a way to defend a strong Indifference Principle from the standard objections.

A and *B* are evidentially symmetric for an agent *iff* she has "no more reason to suppose that *A* is true than that *B* is, or vice versa".[28]

White offers a case for adopting this principle, and a defense of the principle against objections. His arguments in favor of the principle are familiar—these are the considerations discussed in section 3.4. But some of his arguments in defense of the principle are novel, including his response to the cube factory case. In what follows, I'll examine White's response to the cube factory case, and assess how his principle fares with respect to the three worries raised in section 3.4.

4.1 White on the Cube Factory Argument

Let's begin by looking at White's reply to the cube factory case. Let $\approx$ be the evidential symmetry relation. Let L_1/L_2 be the propositions that the cube has a length of (0,1]/(1,2] cm. And let $A_1/A_2/A_3/A_4$ be the propositions that the cube has a face area of (0,1] cm^2/(1,2] cm^2/(2,3] cm^2/(3,4] cm^2. White sets up a Cube Factory Argument against the Indifference Principle as follows:[29]

White's Cube Factory Argument:

(1) $L_1 \approx L_2$ (Premise)

(2) $A_1 \approx A_2 \approx A_3 \approx A_4$ (Premise)

(3) $A \approx B \rightarrow cr(A) = cr(B)$ (White's Indifference Principle)

(4) $cr(L_1) = 1/2$ (1,3)

(5) $cr(A_1) = 1/4$ (2,3)

[28]See White (2010), p.161. White also offers another characterization of evidential symmetry, according to which *A* and *B* are evidentially symmetric "for a subject if his evidence no more supports one than the other" (White (2010), p.161). I've employed White's "reasons" characterization instead of this one for two reasons. First, if we employ the natural Bayesian understanding of "evidential support" (c.f. section 5.1), White's principle becomes vacuous. Second, there's reason to think White is employing a non-standard notion of evidence here (and *a fortiriori*, a non-standard notion of evidential support), making the content of this characterization unclear. White states that "I mean to understand *evidence* very broadly here to encompass whatever we have to go on in forming an opinion about the matter. This can include non-empirical 'evidence', if there is such" (White (2010), p.161-162). If an agent is deciding what her credence in *A* should be, then her certainty that the chance of *A* is 1, her lack of inadmissible evidence with respect to that chance, and the appropriate chance-credence principle are presumably all part of what she "has to go on". Thus her certainty in the chance, her lack of inadmissible evidence, and the chance-credence principle, appear to all (either singly or jointly) count as evidence in White's sense, even though they will not all count as evidence in the standard Bayesian sense.

In any case, which characterization of evidential symmetry we employ has little bearing on what follows. The dialectic proceeds in precisely the same way if we employ this other characterization of evidential symmetry. Just replace all talk of when an agent has no more reason to suppose *A* than *B* with talk of when what an agent "has to go on" no more supports *A* than *B*.

[29]Strictly speaking, this argument also assumes that one's rational credence function *cr* must satisfy the probability axioms, that the L_is and A_is each form a partition of the doxastic possibilities, and that L_1 and A_1 are equivalent propositions, and so must be assigned the same credence. I follow White in leaving these premises implicit, since both sides will grant these assumptions.

(6) $cr(L_1) \neq cr(A_1)$ (4,5)

(7) $cr(L_1) = cr(A_1)$ (Equivalence)

Since (6) and (7) are inconsistent, we have a *reductio* of one of the premises. And since (1) and (2) seem true, this appears to yield a *reductio* of White's Indifference Principle, (3).

White challenges this argument by contesting premises (1) and (2). White argues that we can derive an absurd conclusion from (1) and (2) alone, as follows:[30]

White's Reductio of (1)∧(2):

(1) $L_1 \approx L_2$ (Premise)

(2) $A_1 \approx A_2 \approx A_3 \approx A_4$ (Premise)

(3*) $L_1 \approx A_1$ (Equivalence)

(4*) $L_2 \approx (A_2 \vee A_3 \vee A_4)$ (Equivalence)

(5*) $L_1 \approx (A_2 \vee A_3 \vee A_4)$ (1,4* Transitivity)

(6*) $A_1 \approx (A_2 \vee A_3 \vee A_4)$ (3*,5* Transitivity)

(7*) $A_2 \approx (A_2 \vee A_3 \vee A_4)$ (2,6* Transitivity)

(8*) $A_2 \not\approx (A_2 \vee A_3 \vee A_4)$ (Premise)

Since (7*) and (8*) are inconsistent, we have a *reductio* of one of the premises. White argues that (8*) is true: "surely we have at least some more reason to believe the logically weaker $(A_2 \vee A_3 \vee A_4)$ than to believe A_2."[31] Thus the fault must lie with (1) or (2). At least one of these premises must be false.

Unfortunately, White's reply to the Cube Factory Argument is not compelling. The problem is that the status of White's argument depends on when we take an agent to have "no more reason to suppose A than B". Without some further substantive claims about when one has no more reason to suppose A than B, he cannot show that his reply to the Cube Factory Argument is successful. And White doesn't make any substantive claims about when an agent has no more reason to suppose A than B.

Let's go through this more slowly. First, let's see why White's reply hangs on what further substantive claims one makes about when one has no more reason to suppose A than B. Consider three toy examples of claims one might make about when an agent has no more reason to suppose A than B. (To be clear, these are not supposed to be interpretations of White, nor even plausible proposals; they are just toy examples we are using to show how White's response hangs on what substantive claims we make about reasons.)

Claim 1: One has no more reason to suppose A than B (or vice versa) *iff* the highest credence a rational agent could assign to A is the same as the highest credence she could assign to B.

[30] In addition to the assumptions mentioned in the previous footnote, this argument assumes that L_2 and $A_2 \vee A_3 \vee A_4$ are equivalent propositions, and that the evidential symmetry relation is transitive.

[31] White (2010), p.166.

Claim 2: One has no more reason to suppose A than B (or vice versa) *iff* one ought to assign A and B the same credence.

Claim 3: One has no more reason to suppose A than B (or vice versa) *iff*, for measure μ over the space of possibilities, one's total evidence E is such that $\mu(A|E) = \mu(B|E)$.

Given Claim 1, premise (8*) is false, and White's *reductio* argument against (1)∧(2) fails. A rational agent who gets A_2 as evidence will assign it a credence of 1. So the highest credence a rational agent could assign A_2 is 1. Likewise, the highest credence a rational agent could assign $A_2 \vee A_3 \vee A_4$ is 1. So (8*) is false, since the highest credence a rational agent could assign to A_2 *is* the same as the highest credence a rational agent could assign to $A_2 \vee A_3 \vee A_4$: 1.[32]

On the other hand, given Claims 2 or 3, premise (8*) is true, and White's *reductio* of (1)∧(2) succeeds. Given Claim 2, (1) entails that one's credence in each L_i ought to be the same. Since L_1 and L_2 are mutually exclusive and exhaustive, it follows that one's credence in each should be 1/2. Likewise, (2) entails that one's credence in each A_i ought to be the same. Since $A_1 - A_4$ are mutually exclusive and exhaustive, it follows that one's credence in each should be 1/4. But since A_1 and L_1 are equivalent, they must be assigned the same credence. Thus given Claim 2, we can derive a contradiction from (1) and (2). And similar reasoning shows how to derive a contradiction from (1) and (2) given Claim 3.

So whether White's *reductio* argument works depends on what substantive claims one makes about when an agent has no more reason to suppose A than B. But White doesn't say anything about when an agent has no more reason to suppose A than B. Indeed, as I understand it, White's proposal is a very modest one. White is not trying to make any substantive claims about when an agent has no more reason to suppose A than B, nor is he trying to settle what we ought to believe. He is just proposing a constraint on how our reasons (whatever they are) tie to what we ought to believe (whatever that is).[33]

But by being so modest in his ambitions, White doesn't give himself enough to establish the results he desires. For instance, White can't establish that his principle escapes the Cube Factory Argument. Because one can't show this without making some substantive claims about when an agent has no more reason to suppose A than B. If one doesn't say enough about this to rule out Claim 1, for example, then one leaves open the possibility that the Cube Factory Argument against White's principle succeeds.

Indeed, if this modest understanding of White is correct, it's not clear that White's principle has enough content to bear on the permissive/impermissive debate. Because without making some substantive claims about when an agent has no more reason to suppose A than B, White's principle won't even rule out extreme permissive Bayesianism. If the only substantive constraint we place on reasons is Claim 2, for example, then White's principle tells us that we ought to have the same credence in any two propositions that we ought to have the same

[32]Likewise, as one would expect, premises (1) and (2) are true given Claim 1, and so the Cube Factory Argument against White's Indifference Principle succeeds. The highest credence a rational agent could assign L_1 and L_2 is the same—1—so (1) is true. Likewise, the highest credence a rational agent could assign A_1-A_4 is the same, so (2) is true.

[33]This understanding is suggested by White's comments on p.168 of White (2010), and the passage quoted in footnote 35, below.

credence in. Even the extreme permissive Bayesian can accept that.[34]

4.2 Three Worries

Let's turn to assess how White's proposal fares with respect to the three general worries for Indifference Principles raised in section 3.4.

What about the first worry, that an Indifference Principle faces the trilemma of being either inconsistent, trivial or arbitrary? How White's principle fares with respect to the trilemma will depend on what substantive claims one makes about when an agent has no more reason to suppose A than B.

For example, consider the three toy examples from the last section. Suppose Claim 1 were true—that one has no more reason to suppose A than B *iff* the highest credence a rational agent could assign to A and to B is the same. This entails that $A \approx \neg A \approx A \vee \neg A$. White's principle would then entail that $cr(A) = cr(\neg A) = cr(A \vee \neg A)$, which is probabilistically incoherent. Thus White's principle would fall on the inconsistency horn of the trilemma.

Suppose instead that the only substantive constraint on reasons were provided by Claim 2—that one has no more reason to suppose A than B *iff* one ought to assign A and B the same credence. Then White's principle would boil down to the claim that an agent ought to have the same credence in two propositions *iff* her credence in those propositions ought to be the same. Thus White's principle would fall on the triviality horn of the trilemma.

Suppose Claim 3 were true—that one has no more reason to suppose A than B *iff* one's total evidence E is such that $\mu(A|E) = \mu(B|E)$. Then White's principle would entail that one's initial credences should line up with measure μ. But without some further story about why μ is the correct measure to use, this constraint seems arbitrary. Thus White's principle would threaten to fall on the arbitrariness horn of the trilemma.

More abstractly, all of the choice-of-carving problems that face the standard Indifference Principles manifest themselves as choice-of-reasons problems for White's principle. And just as each way of resolving the choice-of-carving problem leads the standard Indifference Principles to one of the horns of the trilemma, each way of resolving the choice-of-reasons problem leads White's principle to one of the horns of the trilemma. So White's principle appears to fare no better with respect to the trilemma than the standard Indifference Principles.[35]

[34]This point should not be understood as a *criticism* of White, since it's unclear how strong White (2010) takes his Indifference Principle to be. For example, on the modest understanding of White I've suggested, White is officially neutral about whether his principle eliminates all but one priors function (i.e., is a strong Indifference Principle), eliminates all but a restricted set of priors functions, or eliminates no priors functions at all. (Of course, since he is *amenable* to it being a strong Indifference Principle, we need to assess it anyway, in order to find out whether he has found a way to defend (what is potentially) a strong Indifference Principle from the standard objections.) In any case, note that the worries for White's defense discussed in sections 4.1 and 4.2 arise regardless of whether we take it to be a strong Indifference Principle or not.

[35]White briefly discusses the trilemma in the following passage:

> "I suspect that many who are hostile to POI [the Principle of Indifference] view it as trying to do something clearly misguided: taking purely structural features of a space of possibilities as giving conditions on rational credence. The trouble is that there are different structures we can impose on a space. We need something more to tell us which way to cut the pie to get a unique answer. If nothing further is specified our criterion is

What about the second worry, that there's little to be gained by adopting an Indifference Principle? Again, White's proposal doesn't escape this worry. White appeals to indifference intuitions in a number of cases to motivate his principle. But since White's principle doesn't entail that any of these intuitive verdicts are correct, these indifference intuitions don't support his principle. And we still have little reason to believe that these intuitions support a single rationality constraint.

What about the third worry, that the intuitions motivating Indifference Principles are unreliable? Again, White's proposal doesn't escape this worry. Since our indifference intuitions are inconsistent, we have good reason to think they're not trustworthy. And thus we have good reason to disregard the indifference intuitions that White appeals to in order to motivate his principle.

So White's proposal does not yield the Indifference Principle impermissive Bayesians are looking for. In the end, White's principle is subject to the same worries as the other Indifference Principles.

5 White's Case for Evidential Uniqueness

White (2005) presents a series of arguments against permissive views, i.e., views that reject Evidential Uniqueness. In this section, I'll examine four of these arguments.[36]

Three comments before we begin. First, I will assess White's case for Evidential Uniqueness in the context of Bayesianism. If White's arguments work against permissivist views in general, then they should work just as well against particular permissivist views, like permissive Bayesianismism. And if these arguments fail to rule out permissive Bayesianism, then they fail as arguments against permissivism full stop. In any case, White himself is a Bayesian, and has the Bayesian case in mind when presenting these arguments.[37]

Second, some of White's arguments presuppose that permissivists reject Agent Uniqueness, while others presuppose that permissivists accept Permission Parity.[38] As we saw in sec-

> empty. If all carvings are allowed we get inconsistency. If further criteria are imposed they often seem arbitrary or unmotivated. ... POI as I'm understanding it is importantly different. It takes an *epistemic* input ("having no more reason...") to deliver an epistemic output (equal credence). This is not open to the same charge of arbitrariness. It is appropriate that facts about the balance of my reasons should put constraints on my credal states." (White (2010), p.168.)

But again, this claim does not seem sufficient. Without making substantive claims about when an agent has no more reason to suppose A than B, one can't show that the principle avoids the inconsistency horn, the triviality horn, or the arbitrariness horn.

[36] See Kelly (2012) for a discussion of some of the other arguments that White offers.

[37] See White (2005), p.458, footnote 12.

[38] For example, one version of the first premise of the Evidential Support Argument (P1′ of section 5.1), the first premise of the Flip-Flopping Argument (P1 of section 5.2), the first premise of the Truth-Guiding Argument (P1 of section 5.3), and the first premise of the Practical Deliberation Argument (P1 of section 5.4), would be false if not restricted to some version of permissive Bayesianism that rejects Agent Uniqueness. And the second premise of the Truth-Guiding Argument (P2 of section 5.3) must be restricted to views which accept Permission Parity in order to not be subject to obvious counterexamples.

tion 2, this need not be the case: those who reject Evidential Uniqueness need not reject Agent Uniqueness or accept Permission Parity. But I think White's arguments are problematic for independent reasons. So for the sake of argument, I will assume in what follows that the only viable form of permissive Bayesianism is permissive$_1$ Bayesianism. (Permissive$_1$ Bayesians, recall, are Bayesians who reject Agent Uniqueness and accept Permission Parity, holding that there are multiple permissible priors functions for each agent, and that which priors functions are permissible is the same for all possible agents.)

Third, some of White's arguments are presented as targeting only extreme permissivism, while others are presented as targeting any form of permissivism. But we needn't concern ourselves with this distinction here. White raises the same kinds of worries for both moderate and extreme forms of permissivism. And the arguments we'll look at will apply to any kind of permissivism.

5.1 The Evidential Support Argument

First, let's look at what I'll call the "Evidential Support Argument".

> "How could evaluation of the evidence render me rational in believing that Smith is guilty...? Surely only if the total evidence supports Smith's guilt. ... But the evidence cannot support both Smith's innocence and his guilt. ... it is impossible that my examination of the evidence makes it rational for me to believe that Smith is guilty but also rational to believe instead that he is innocent. ... [This] departure from Uniqueness cannot be right."[39]

Here is an argument one might construct from this passage:

The Evidential Support Argument:

P1. If permissive$_1$ Bayesianism is true, then some total evidence E could both support A and support $\neg A$.

P2. No total evidence E could both support A and support $\neg A$.

C. Permissive$_1$ Bayesianism is false.

Before we can assess this argument, we need to figure out what it means to say that total evidence E supports A. Presumably, if total evidence E supports A, then a rational agent with no evidence who gets E should increase her credence in A. But whether such an agent should increase her credence in A after getting E will depend on her background beliefs.

For example, suppose a rational agent is uncertain about whether a certain machine is on. If she antecedently believes that the red light on the machine turns on when the machine is working, then seeing that the red light is on should increase her credence that the machine is working. On the other hand, if she antecedently believes that the red light turns on when the machine is broken, then seeing that the red light is on should decrease her credence that the machine is working.

[39] White (2005), p.447.

More generally, anyone who accepts conditionalization will take the evidential support relation to be a three place relation, since three things are relevant to what one's credence should be: the proposition, the evidence, and one's background beliefs. And if we're considering what's supported by one's *total* evidence, these relata will be: the proposition, one's total evidence, and one's priors. This suggests the following characterization of total evidential support:

Total Evidential Support: Total evidence E supports A with respect to rational priors function ic *iff* $ic(A|E) > ic(A)$.

Impermissive Bayesians hold that there's only one rational priors function. So impermissive Bayesians can leave the priors argument of the total evidential support relation implicit. But since we want to set up the Evidential Support Argument in a way that's neutral between permissive and impermissive Bayesianism, we have to keep the priors argument explicit.[40]

Now let's turn to assess the Argument from Evidential Support. Since "total evidential support" is a three-place relation, the first and second premises of the argument are ambiguous. First, we might understand these premises as follows:

P1.′ If permissive$_1$ Bayesianism is true, then some total evidence E could support A with respect to one rational priors function and support $\neg A$ with respect to another.

P2.′ No total evidence E could support A with respect to one rational priors function and support $\neg A$ with respect to another.

On this understanding of the argument, P1′ is true. But P2′ is question begging, since (assuming Bayesianism) it is just a re-statement of Evidential Uniqueness. So this version of the argument assumes what it is trying to show.[41]

Second, we might understand the first and second premises of the argument as follows:

P1.″ If permissive$_1$ Bayesianism is true, then some total evidence E could both support A and $\neg A$ with respect to some particular rational ic.

P2.″ No total evidence E could both support A and $\neg A$ with respect to some particular rational ic.

On this understanding of the argument, P2″ is not contentious. But P1″ is false: permissive Bayesians$_1$ will not hold that one's total evidence E could both support A and $\neg A$ with respect to some particular ic.

Let's consider one more possibility. Perhaps "evidential support" should be understood as a primitive, intuitive notion, not one that can be analyzed in Bayesian terms in the manner I sug-

[40] Kelly (2012) raises similar worries.

[41] Christensen (2007) seems to offer a similar argument: "There's something unstable about holding onto my belief while acknowledging that a different belief enjoys equal support from the evidence." (Christensen (2007), p.191.) Why unstable? Well, if we assume that the total evidential support relation has only two changeable arguments—one's total evidence and the proposition in question—then either my total evidence supports A or it doesn't. And if I believe A while believing that my total evidence supports $\neg A$, then I don't believe what I think I should. In this sense, my beliefs are unstable. The problem with this line of thought is that by assuming the priors argument of the total evidential support relation is fixed, it effectively assumes Evidential Uniqueness from the outset.

gest above.[42] But on this approach, unless more is said, we have no reason to accept P1—that one's total evidence could support both A and $\neg A$ given permissive Bayesianism. Permissive Bayesianism doesn't make any claims about a primitive notion of evidential support. To make a claim like P1, we need to have some way of tying this primitive notion of evidential support to what permissive Bayesianism actually makes claims about—what one's credences ought to be. And once we add constraints tying the evidential support relation to what one's credences ought to be, the dialectic reverts to the one we have just been through.

So the Evidential Support Argument is either question begging or unsound. Either way, it fails to provide us with a reason to adopt Evidential Uniqueness.

5.2 The Flip-Flopping Argument

Next, let's look at what I'll call the "Flip-Flopping Argument".

> "[I]f I really do judge that believing P in this situation would be rational, as would believing $\neg P$, then there should be nothing wrong with my bringing it about that I have some belief or other on the matter. But then it surely cannot matter how I go about choosing which belief to hold, whether by choosing a belief that I'd like to hold, or flipping a coin, or whatever. ... [Adopting beliefs this way would be irrational, so] we have reached the conclusion that I cannot rationally accept the extreme permissivist thesis..."[43]

Here is one argument that might be drawn from this passage:[44]

The Flip-Flopping Argument (v1):

P1. If permissive$_1$ Bayesianism is true, then it can be permissible in a particular situation to have a credence of either x or y in A (where $x \neq y$).

P2. If it is permissible in a particular situation to have a credence of either x or y in A, then there is nothing wrong with flip-flopping on A—i.e., arbitrarily changing one's credence in A from x to y or vice versa.

P3. There is something wrong with arbitrarily changing one's credence in A from x to y or vice versa.

C. Permissive$_1$ Bayesianism is false.

---

[42] I thank a referee for suggesting this understanding.

[43] White (2005), p.449-450. Similar considerations are raised in other parts of the paper, such as p.453 and p.454-455.

[44] The argument presented here is somewhat simpler than the argument suggested in this passage. Unlike the argument suggested in the passage, I've formulated the argument in terms of what it's rational to believe, not what one *believes* it's rational to believe. One could provide a more accurate reconstruction of White's argument by formulating the argument in terms of what one believes to be permissible, and adding the premise that in order for an account of epistemic norms to be correct, it must be permissible to believe it. But this merely serves to complicate the argument, and this more complex version of the argument runs into the same difficulties as the simpler argument presented in the text.

The first premise of this argument is true. In most cases, permissive Bayesians$_1$ will maintain that an agent's beliefs are fixed by her evidence and her prior beliefs. But in *tabula rasa* cases, a number of different doxastic states are permissible, since an agent's evidence and (non-existent) prior beliefs won't fix what she ought to believe. Thus, given permissive$_1$ Bayesianism, there are situations in which it can be permissible to have a credence of either x or y in A.

No Bayesian will accept the second premise, however. Arbitrarily changing one's beliefs without getting new evidence violates conditionalization. Since Bayesians accept conditionalization, they will reject the second premise.

If P2 is false, why does it initially appear to be plausible? One possibility is that it's easy to mistake P2 for the following claim:

P2.′ If it is permissible in a particular situation to have a credence of either x or y in A, then there is nothing wrong with having a credence of either x or y in A in this situation.

P2′ is clearly true. But despite their resemblance, P2 and P2′ are different claims. While P2′ states a tautology, P2 makes a substantive deontic claim.

With this in mind, we might draw a different argument from the passage above, an argument which employs P2′ and P3′ instead of P2 and P3:

The Flip-Flopping Argument (v2):

P1. If permissive$_1$ Bayesianism is true, then it can be permissible in a particular situation to have a credence of either x or y in A (where $x \neq y$).

P2.′ If it is permissible in a particular situation to have a credence of either x or y in A, then there is nothing wrong with having a credence of either x or y in A in this situation.

P3.′ There is something wrong with having a credence of either x or y in A in this situation.

C. Permissive$_1$ Bayesianism is false.

Now the second premise of the argument is unproblematic. But the third premise of the argument is question begging, since it presupposes Agent Uniqueness, which the permissive$_1$ Bayesian rejects.[45]

Here is a third way of drawing an argument from the passage above, which focuses on the process by which one forms a belief:[46]

The Flip-Flopping Argument (v3):

P1. If permissive$_1$ Bayesianism is true, then it can be permissible in a particular situation to have a credence of either x or y in A (where $x \neq y$).

P2.″ If it is permissible in a particular situation to have a credence of either x or y in A, then there is nothing wrong with (say) flipping a coin to determine which credence to adopt.

P3.″ There is something wrong with flipping a coin to determine which credence to adopt.

[45] What if we instead considered permissive Bayesians who do accept Agent Uniqueness, like permissive$_3$ Bayesians? Then P1 would be false, so the argument would be unsound.

[46] Thanks to Jonathan Weisberg here for suggesting this understanding.

C. Permissive$_1$ Bayesianism is false.

This argument is unsound because the second premise is false. We can see this by constructing a counterexample to the second premise.

The core Bayesian principles, Probabilism and Conditionalization, are both silent about the process by which one comes to have a belief. (Conditionalization will impose constraints on how one's beliefs at different times must be related, of course. But it doesn't impose constraints on the process by which one comes to have these beliefs. As long as these beliefs line up in the right way, Conditionalization doesn't care whether these beliefs were formed by rational deliberation, hypnosis, quantum mechanical fluctuations, or what have you.) While Probabilism and Conditionalization don't impose constraints on how one comes to have a belief, they don't forbid such constraints either. So a Bayesian can hold both that it is permissible to have a credence of either x or y in A, and that it's only permissible to have these credences if one comes to have them via the right kind of belief-forming process.[47] Thus the conditional that the second premise asserts is false.

Let's step back and consider this kind of argument from a more general perspective. The problem with this kind of argument is that it tries to settle the Evidential Uniqueness debate by appealing to considerations regarding the process by which one comes to have a belief. But these kinds of "process" considerations are orthogonal to the Evidential Uniqueness debate.

Consider the following two claims:

Process: There are normative constraints on the process by which one comes to have a belief.

No Process: There are no normative constraints on the process by which one comes to have a belief.

Evidential Uniqueness is compatible with both Process and No Process. One can accept Evidential Uniqueness and hold that there are no constraints on how one comes to have one's beliefs—all that matters is that one's beliefs that line up correctly with one's evidence. Whether one came to have those beliefs as the result of diligent study, coin-flipping, or brain washing, is irrelevant. Likewise, one can accept Evidential Uniqueness and hold that there *are* constraints on how one comes to have one's beliefs—one should not only have beliefs that line up correctly with one's evidence, one should also have arrived at those beliefs in the right way.

Similarly, the denial of Evidential Uniqueness is compatible with both Process and No Process. One can reject Evidential Uniqueness and hold that there are no constraints on how one comes to have one's beliefs—all that matters is (say) that one comes to have the beliefs that line up in the right way with one's evidence, one's prior beliefs, and whatever else one takes to be relevant. Likewise, one can reject Evidential Uniqueness and hold that there *are* constraints on how one comes to have one's beliefs—not only should one's beliefs line up with one's evidence, one's prior beliefs, etc., but one should also come to have those beliefs in the right way.

Thus the Process/No Process debate is independent of the Evidential Uniqueness debate. The Evidential Uniqueness debate is about what normative constraints one's evidence at a time

[47] E.g., a permissive Bayesian might hold that it's permissible to have a credence of either x or y in A in a *tabula rasa* case, but also hold that these initial credences are only permissible if one came to have one's initial credences via the right kind of process.

places on one's doxastic state at that time. Whether there are also normative constraints on how one comes to have that doxastic state is orthogonal to this debate.

Let's sum up. On the first way of understanding the Flip-Flopping Argument it conflicts with Bayesianism. On the second way of understanding the argument it's question begging. On the third way of understanding the argument it's unsound. On any of these understandings, the Flip-Flopping Argument fails to provide us with a reason to adopt Evidential Uniqueness.

5.3 The Truth-Guiding Argument

Now let's look at what I take to be White's main argument, which I'll call the "Truth-Guiding Argument".

> "Now let's suppose... that upon considering the evidence in court one could rationally conclude that Smith is guilty, but there is an alternative path that one's reasoning could take arriving instead at the rational conclusion that he is innocent. Supposing this is so, is there any advantage, from the point of view of pursuing the truth, in carefully weighing evidence to draw a conclusion, rather than just taking a belief-inducing pill? Surely I have no better chance of forming a true belief either way."[48]

Here is an argument one might extract from this passage:

The Truth-Guiding Argument

P1. If permissive$_1$ Bayesianism is true, there are some situations in which multiple doxastic states are permissible.

P2. If an account of epistemic norms satisfies Permission Parity (as permissive$_1$ Bayesianism does), and takes there to be some situations in which multiple doxastic states are permissible, then it cannot be truth-guiding.

P3. The correct account of epistemic norms is truth-guiding.

C. Permissive$_1$ Bayesianism is false.

First, a quick note about the second premise. The second premise is restricted to accounts which satisfy Permission Parity in order to avoid obvious counterexamples. Accounts which satisfy Permission Parity take the same priors to be permissible at all worlds. Thus the priors they permit can't vary to take what the world is like into account. But accounts that reject Permission Parity can maintain that which priors are permissible depends on what the world is like. And such views can be as truth-guiding as one likes. Thus if we don't restrict the second premise to accounts which satisfy Permission Parity, it will be subject to straightforward counterexamples. For example, consider a view which requires agents to have priors that assign a credence greater than 0.99 to every true proposition. In *tabula rasa* cases, this view permits multiple doxastic states. Yet it will clearly be truth-guiding, given any reasonable notion of "truth-guiding".

[48] White (2005), p.448. White raises similar considerations in other parts of the paper, such as p.449, p.451-452, and p.457.

Now let's turn to consider the Truth-Guiding Argument as a whole. The first premise of this argument is a restatement of the first premise of the Flip-Flopping Argument, and is true. So the status of this argument hangs on the viability of the second and third premises. And in particular, it hangs on the notion of truth-guiding that these premises employ. Since there's a wide range of different notions of truth-guiding that one could employ, it's hard to say anything comprehensive about this argument. Nevertheless, there are reasons to think that this argument is not compelling.

To begin, there are reasons why Bayesians of any stripe might be wary of truth-guiding requirements. While most Bayesians are concerned with truth, they are also concerned with other desiderata, desiderata which are orthogonal to or in conflict with some notions of truth-guiding. And they might take certain truth-guiding desiderata to be outweighed by these other considerations.

For example, a number of Bayesians have appealed to Dutch book arguments to justify Bayesian norms. These arguments claim that if one violates Bayesian norms, then one can be led to accept bets which together ensure a loss. Susceptibility to Dutch books might be taken to indicate epistemic incoherence, and thus to provide us with epistemic reasons to be Bayesian.[49] And an epistemic norm which fails these coherence desiderata but satisfies certain truth-guiding requirements might be less appealing, all things considered, than a rival that maintains coherence but does not satisfy these requirements.

Here is another example, drawn from recent discussions by Dunn (2011) and Hajek (2011). Consider some time-dependent chancy event, say the outcome of a coin toss.[50] Suppose the chance of this coin landings heads is currently 0.5. And suppose that, as it turns out, the coin will, in fact, land heads. Is it better to have a credence of 0.5 that the coin will land heads, or a credence of 1 that the coin will land heads?[51] If one adopts a "calibration" perspective, according to which our epistemic aim should be to have credences that line up with the current chances, it is epistemically better to have a credence of 0.5 that the coin will land heads. If one adopts an "accuracy" perspective, according to which our epistemic aim should be to have credences that are close to the truth as possible, it is epistemically better to have a credence of 1 that the coin will land heads. The truth-guiding requirement naturally suggests something like the accuracy perspective. But, as this example shows, this demand for truth can conflict with the demand for calibration. And some Bayesians may maintain that when these two desiderata conflict, the desiderata of calibration should trump that of truth.

[49] See Christensen (1996).

[50] Of course, it's contentious whether ordinary coin tosses are chancy (see Lewis (1986)). And among those who think they are chancy, it is contentious whether these chances should be thought of as time-dependent (see Meacham (2005), Hoefer (2007)). To get around these worries, we could replace the coin toss with the spin measurement of an electron, given an indeterministic interpretation of quantum mechanics like GRW. Likewise, one could restrict the scope of the calibration perspective described below to indeterministic time-dependent chances. (Alternatively, one could generalize the calibration perspective by formulating it in terms of something like conditional chances.)

[51] Note that this question is independent of what evidence one has. We're considering what epistemic states it's epistemically better to be in given what the world is actually like, full stop. (Of course, one might flat out reject this kind of "epistemic axiology". If so, one would be unhappy with both the calibration and the accuracy perspectives described below.)

Here is a third example. It is sometimes held that a desiderata of an account of epistemic norms is that it provide us with useful advice. Consider an account of epistemic norms which states only that one should adopt all and only true beliefs. This account may do well according to certain truth-guiding requirements, but it offers little useful epistemic advice.[52] As this example shows, the demand for truth can conflict with the demand for usefulness. And in cases of conflict, some will take the desiderata of usefulness to trump that of truth.

Thus there are several reasons why Bayesians might be wary of truth-guiding requirements. Now let's turn to some other concerns.

The Truth-Guiding Argument attempts to establish that permissive$_1$ Bayesianism fails to be appropriately truth-guiding. For this argument to be compelling, we must employ a notion of truth-guiding that it's reasonable to expect a theory to satisfy. We should only expect permissive$_1$ Bayesianism to be truth-guiding in ways that we would expect its impermissive rivals to be truth-guiding. Thus it's natural to evaluate the Truth-Guiding Argument in comparative terms: is permissive$_1$ Bayesianism as truth-guiding as impermissive Bayesianism?

Given some notions of truth-guiding, both permissive$_1$ and impermissive Bayesianism will be truth-guiding.[53] For example, there are a number of convergence theorems in the literature that show (roughly) that as Bayesian agents accumulate evidence, there's a high probability that their credence in true propositions will increase.[54] Of course, these results come with a number of caveats: some of them only apply in the long run, some of them only show that an agent will *believe* that her credence in true propositions will increase, and all of these results require substantive assumptions about an agent's evidence, such as that this evidence is veridical and sufficiently discriminating, or believed by the agent to be so. Nevertheless, given a suitably relaxed understanding of truth-guiding, these results demonstrate some ways in which all forms of Bayesianism are truth-guiding.

Given other notions of truth-guiding, neither permissive$_1$ nor impermissive Bayesianism will be truth-guiding.[55] Consider the worry raised by White's remarks from above. If "one could rationally conclude that Smith is guilty, but there is an alternative path that one's reasoning could take arriving instead at the rational conclusion that he is innocent", then there is no "advantage, from the point of view of pursuing the truth, in carefully weighing evidence to draw a conclusion, rather than just taking a belief-inducing pill".[56] Here is one way to flesh out this worry. Consider a permissive Bayesian who takes both ic and ic^* to be rationally permissible priors. Given evidence E, the permissive Bayesian may tell an agent with priors ic to have a high credence in A, and an agent with priors ic^* to have a low credence in A. But, of course, either A is true or it's not. So one of these agents will have a credence in A that is far from the truth, even though both agents conform perfectly to the norms of rationality. Thus, if

[52] Of course, Bayesians generally employ a highly idealized conception of the agents in question. So this desiderata might motivate other changes to the standard Bayesian account as well.

[53] Given these notions of truth-guiding, the argument is unsound because P2—that a Permission Parity-satisfying account of epistemic norms cannot be truth-guiding if it takes it to be permissible in some situations to have any one of several doxastic states—is false.

[54] See Hawthorne (2012), and the references therein.

[55] Given these notions of truth-guiding, the argument seems unsound because P3—that the correct account of epistemic norms is truth-guiding—seems false.

[56] White (2005), p.448.

permissive$_1$ Bayesianism is true, it follows that an agent can conform perfectly to the norms of rationality, and still end up with beliefs that are far from the truth.

This is true, of course. But impermissive Bayesians are no better off. Consider an impermissive Bayesian who takes *ic* to be the only rationally permissible prior. Given evidence E, the impermissive Bayesian will tell an agent to have a high credence in A. But one's total evidence doesn't generally fix what the world is like.[57] One can get evidence E at a world where A is true or a world where $\neg A$ is true. So, at some worlds, the impermissive Bayesian will tell an agent to have a high credence in A, even though A is false. Thus, if impermissive Bayesianism is true, it follows that an agent can conform perfectly to the norms of rationality, and still end up with beliefs that are far from the truth.[58]

So if we adopt the notion of truth-guiding suggested by this worry—where an account of norms is only truth-guiding if an agent who conforms perfectly to those norms won't end up with beliefs far from the truth—then neither permissive$_1$ nor impermissive Bayesianism will qualify as truth-guiding.

Finally, there are tailored notions of truth-guiding that do discriminate between impermissive and permissive$_1$ Bayesianism. But these notions of truth-guiding do little to bolster the Truth-Guiding Argument. For while one can tailor notions of truth-guiding to favor impermissive over permissive$_1$ Bayesianism, one can also tailor notions of truth-guiding to favor permissive$_1$ over impermissive Bayesianism. And the notions of truth-guiding that favor impermissive Bayesianism are no more plausible than those that favor permissive$_1$ Bayesianism.

To see this, let's look more carefully at one way of assessing how truth-guiding an account of epistemic norms is. The basic idea is to work out the prescriptions of these norms for the agents at each world, see how close these prescriptions are to what is true at that world, and then use these results to form an overall assessment of how truth-guiding these norms are. By fleshing out these details in different ways, we can get different notions of truth-guiding, some of which favor impermissive Bayesianism, and some of which favor permissive$_1$ Bayesianism.

For example, one detail that needs to be fleshed out is how to compare the value of making it obligatory to have a true belief with the value of making it permissible to have a true belief. Here is one way to assess the value of such prescriptions. Take all of the beliefs the norms permit you to have, and assess these norms using the permitted beliefs that are farthest from the truth. This way of assessing the value of prescriptions will favor impermissive Bayesianism, since adding additional permissible priors can only hurt an account of epistemic norms.

Here is another way to assess the value of such prescriptions. Take all of the beliefs the norms permit you to have, and assess these norms using the permitted beliefs that are closest to the truth. This way of assessing the value of prescriptions will favor permissive$_1$ Bayesianism, since adding additional permissible priors can only help an account of epistemic norms.

[57] One might come up with cases in which one's evidence does fix what the world is like—one might be a deity, say, who receives as veridical evidence the conjunction of every true proposition. But in these cases, Bayesians of any stripe will prescribe true beliefs, so the impermissive Bayesian gains no ground on the permissive$_1$ Bayesian.

[58] North (2010) raises similar worries.

Another detail that needs to be fleshed out is how to weigh the value of prescribing a true belief at one world versus another.[59] For example, one might want to employ a measure which assigns a smaller weight to "misleading" worlds than to "non-misleading" worlds. And there are an infinite number of different measures to choose from for weighing the value of true beliefs at different worlds.

Priors that assign higher values to worlds with greater weight will generally do better with respect to truth-guiding than priors that don't. So weighing measures will generally favor priors that line up with them. With this in mind, one could argue that truth-guiding considerations favor a given form of impermissive Bayesianism, by maintaining that there's a unique correct weighing measure, and that a particular initial credence function does better with respect to this measure than any other initial credence function. Likewise, one could argue that truth-guiding considerations suggest a given form of permissive$_1$ Bayesianism, by maintaining that there are several equally good weighing measures, and that all and only the members of some set of initial credence functions do best with respect to these measures.

Let's take stock. According to some notions of truth-guiding both impermissive and permissive$_1$ Bayesianism are truth-guiding, according to others neither are truth-guiding, and according to yet others only impermissive or permissive$_1$ Bayesianism is truth-guiding. In order for the Truth-Guiding Argument to be compelling, the normatively relevant notion of truth-guiding must be one according to which only impermissive Bayesianism is truth-guiding. But the claim that only impermissive Bayesianism is truth-guiding is no more plausible than the claim that only permissive$_1$ Bayesianism is. So truth-guiding considerations don't give us a reason to favor impermissive Bayesianism over permissive$_1$ Bayesianism.

The Truth-Guiding Argument was formulated as an argument against permissive$_1$ Bayesianism. And, so understood, I've suggested that there are reasons to think the Truth-Guiding Argument is not compelling. But we can also consider more general versions of the Truth-Guiding Argument, where the target is any form of permissive Bayesianism, or even any theory that rejects Evidential Uniqueness. And with respect to these more general versions of the Truth-Guiding Argument, we can say something stronger—given any reasonable notion of truth-guiding, these arguments are demonstrably unsound.

One problem is that the analog of P1—that if one of these views is true, there are situations in which multiple doxastic states are possible—will no longer hold. Some of these permissive views will satisfy Agent Uniqueness, and so won't allow for situations in which multiple doxastic states are permissible. Another problem is that we can no longer use P2—that if a Permission Parity-satisfying account allows for situations in which multiple doxastic states are permissible, then it can't be truth-guiding—to establish that all of these views aren't truth-guiding. Since some of these permissive views don't satisfy Permission Parity, P2 won't apply to them. And if we remove the Permission Parity clause from P2 in order to get it to apply to these views, then the premise becomes implausible, for the reasons noted earlier.

Indeed, once we expand our attention to all permissive views, we can see that truth-guiding considerations not only fail to tell against permissive views, they give us a reason to favor permissive views over impermissive views. Theories which accept Evidential Uniqueness are

[59] A similar issue arises with respect to how to weigh the value of prescribing a true belief to one agent at a world versus another.

necessarily insensitive to what the world is like in a certain respect. Since the only factor they can take into consideration is one's evidence, they have to prescribe the same beliefs in evidentially identical situations. And when different things are true in evidentially identical situations, they can't consistently prescribe true beliefs to both. Theories that reject Evidential Uniqueness, on the other hand, are not bound by this restriction. They can take anything into consideration—including what's true!—and so can be as truth-guiding as you like.

It's worth taking a step back to consider these kinds of truth-guiding considerations from a more general perspective. In broad strokes, something like the following thought might motivate one to adopt Evidential Uniqueness:

> A theory that rejects Evidential Uniqueness is no more likely to lead one to true beliefs than (say) flipping a coin. But the right epistemic theory must be more likely to lead one to true beliefs than that. So the right epistemic theory must accept Evidential Uniqueness.

The three concerns that I've raised regarding the Truth-Guiding Argument will apply to any argument along these lines. First, there are reasons why one might be wary of this kind of truth-guiding requirement. Second, restricting our attention to views like impermissive and permissive$_1$ Bayesianism, this kind of argument does not support the former over the latter, since impermissive Bayesianism is no better than permissive$_1$ Bayesianism at leading agents to true beliefs. Third, if we extend our attention to *all* views that reject Evidential Uniqueness, we can see that (a) there are theories that reject Evidential Uniqueness which are as truth-guiding as you like, so the claim that "a theory that rejects Evidential Uniqueness is no more likely to lead one to true beliefs than (say) flipping a coin" is false, and (b) there are theories that reject Evidential Uniqueness that will do better at guiding subjects to true beliefs than any theory that accepts Evidential Uniqueness.

Let's sum up. As originally formulated—as an argument against permissive Bayesianism$_1$—the Truth-Guiding Argument is not compelling. And once we extend the argument to any form of permissive Bayesianism, or to any theory that rejects Evidential Uniqueness, the Truth-Guiding Argument becomes demonstrably unsound. So the Truth-Guiding Argument fails to provide us with a reason to adopt Evidential Uniqueness.

5.4 The Practical Deliberation Argument

Finally, let's look at what I'll call the "Practical Deliberation Argument".

> "Suppose that having carefully considered the evidence, my conviction in Smith's guilt is strong enough to warrant voting Guilty. Surely this is what I should do—to let him go free given that I am (rationally) confident that he's guilty would be grossly irresponsible. ... suppose that the range of rationally permissible degrees of confidence in Smith's guilt is just broad enough to include my own conviction as well as a degree of doubt that would make a vote of Not Guilty appropriate. It appears that if this were so I should have no qualms about letting Smith go free, even though I'm sure he's a murderer. ... As long as my vote corresponds to an attitude that is within the bounds of rationally permissible conviction, I have done

> as well as I possibly can given the available evidence to attain a just verdict. ... But of course it can't be right to try to let someone go free when you are rationally persuaded that he is a murderer."[60]

So far, we've been focused on questions regarding what it's permissible to believe, i.e., what's epistemically permissible. This argument also raises questions regarding what it's permissible to do, i.e., what's prudentially permissible. So in what follows, I'll add subscripts to distinguish between talk of what's epistemic permissible (or permissible$_e$) and what's prudential permissible (permissible$_p$).

Here is an argument one might construct from this passage:

The Practical Deliberation Argument:

P1. If permissive$_1$ Bayesianism is true, then the following conditions can jointly obtain:

(i) an agent has credences cr,

(ii) it's permissible$_e$ for her to have credences cr,

(iii) it's permissible$_e$ for her to have credences cr^*.

P2. If conditions (i) and (ii) obtain (an agent has credences cr and it's permissible$_e$ for her to have those credences), then it's impermissible$_p$ for her to perform an act which fails to maximize expected utility with respect to cr.

P3. If conditions (ii) and (iii) obtain (it's permissible$_e$ for an agent to have both credences cr and cr^*), then it's permissible$_p$ for her to perform an act which fails to maximize expected utility with respect to cr.

4. Conditions (i)-(iii) can't jointly obtain (by P2+P3).

C. Permissive$_1$ Bayesianism is false.

The first premise of the argument is plausible, for much the same reasons as the first premise of the Flip-Flopping Argument. But the second and third premises of the argument are jointly problematic. The status of these premises hangs on our understanding of prudential obligation. But no understanding of prudential obligation will make both P2 and P3 true.

The canonical way of thinking about prudential obligation—call it the *subjective* way—takes an act to be permissible$_p$ *iff* it maximizes expected utility with respect to one's credences. Given this subjective understanding of prudential obligation, P2 is true. If one's credences are cr, then it's impermissible$_p$ to perform an act that fails to maximize expected utility with respect to cr. But this subjective understanding of prudential obligation makes P3 false. Even if it's permissible$_e$ for an agent to have either credences cr or cr^*, it doesn't follow that it's permissible$_p$ for her to act in accordance with either one of them. It's only permissible$_p$ for her

[60] White (2010), p.453-454. White raises similar considerations in section 6 of his paper. (The version presented in section 6 is cashed out in terms of what one *believes* it's rational to believe, not what it's actually rational to believe, and has the conclusion that one shouldn't believe impermissivism is true, not that impermissivism isn't true. One can turn this into an argument against impermissivism by adding the premise that it must be permissible to believe the correct account of epistemic norms. In any case, this complication doesn't change the status of the argument; both versions run into the same kinds of difficulties.)

to act in accordance with the credences she actually has. And if she actually has credences cr, then it's not going to be permissible$_p$ for her to perform an act that fails to maximize expected utility with respect to cr.

Another way of thinking about prudential obligation—call it the *prospective* way—takes an act to be permissible$_p$ *iff* it maximizes expected utility with respect to some permissible$_e$ credence function.[61] Given this prospective understanding of prudential obligation, P3 is true. Since cr^* is permissible$_e$, it's permissible$_p$ to perform an act that maximizes expected utility with respect to cr^* but not cr. But this prospective understanding of prudential obligation makes P2 false. It *can* be permissible$_p$ for an agent with permissible$_e$ credences cr to perform an act that fails to maximize expected utility with respect to cr, as long as it maximizes expected utility with respect to cr^*.

One can construct a number of other notions of prudential obligation, but they all encounter similar problems.[62] These difficulties stem from the fact that the argument relies on two conflicting intuitions. Supporting P2 is the intuition that you shouldn't do anything you rationally believe you shouldn't do. Supporting P3 is the intuition that it's permissible to do anything you could rationally believe it's permissible to do. But these intuitions are only plausible given different notions of prudential obligation. The first intuition requires a notion of prudential obligation which allows the rational credences one holds to trump those one doesn't hold. Otherwise the fact that you rationally hold credences cr wouldn't be enough to rule out the permissibility$_p$ of acting in accordance with some other credence function cr^*. The second intuition requires a notion of prudential obligation which doesn't allow the rational credences one actually holds to trump those one doesn't hold. Otherwise, the mere rationality of some credence function wouldn't be enough to establish that it's permissible$_p$ to act in accordance with it. And no notion of prudential obligation can satisfy both of these intuitions.

Our understanding of the Practical Deliberation Argument hangs on the notion of prudential obligation that we employ. But on any understanding of prudential obligation, the argument is unsound. So the Practical Deliberation Argument fails to provide us with a reason to adopt Evidential Uniqueness.

6 Conclusion

These investigations have suggested a push towards the middle. Extreme permissive Bayesianism is difficult to maintain, since it's incompatible with widely accepted norms like Chance-Credence Principles. Impermissive Bayesianism also faces a number of challenges. Many have

[61] See Zimmerman (2006).

[62] For example, one might consider the *objective* notion of obligation, which makes an act permissible$_p$ *iff* it maximizes utility. Or one might consider a notion of prudential obligation which makes an act permissible$_p$ *iff* it maximizes expected utility with respect to either one's actual credences or some permissible$_e$ credences. Or one might consider a notion of prudential obligation which makes it's obligatory$_p$ to perform an act that maximizes expected utility with respect to one's actual credences if they're permissible$_e$; and otherwise makes it permissible$_p$ to perform an act which maximizes expected utility with respect to any permissible$_e$ credence function. But none of these notions will make both P2 and P3 true. (The first makes both P2 and P3 false, the second makes P2 false, and the third makes P3 false.)

thought that in order to obtain the strong constraints on belief that impermissive Bayesianism requires, one must appeal to something like an Indifference Principle, but these principles have poor prospects.[63] And the positive arguments that have been offered in favor of impermissive Bayesianism are not compelling. In light of these difficulties, moderate permissive Bayesianism appears to be in the best shape.[64]

References

Christensen, David. 1996. "Dutch Books Depragmatized: Epistemic Consistency for Partial Believers." *Journal of Philosophy* 93:450–479.

Christensen, David. 2007. "Epistemology of Disagreement: the Good News." *The Philosophical Review* 116:187–217.

Dunn, Jeff. 2011. "Reliability for Degrees of Belief." Unpublished Manuscript.

Feldman, Richard. 2007. Reasonable Religious Disagreements. In *Philosophers Without Gods: Meditations on Atheism and the Secular*, ed. Louise Antony. Oxford University Press pp. 194–214.

Gaifman, Haim. 1988. A Theory of Higher Order Probabilities. In *Causation, Chance and Credence*, ed. Brian Skyrms and William Harper. Kluwer pp. 191–219.

Goodman, Nelson. 1954. *Fact, Fiction and Forecast*. Harvard University Press.

Hajek, Alan. 2011. "A Puzzle About Degree of Belief." Unpublished Manuscript.

Hall, Ned. 1994. "Correcting the Guide to Objective Chance." *Mind* 103:505–517.

Hawthorne, James. 2012. "Inductive Logic." Stanford Encyclopedia of Philosophy, Edward N. Zalta (ed.). http://plato.stanford.edu/archives/win2012/entries/logic-inductive/.

Hoefer, Carl. 2007. "The Third Way on Objective Probability: A Skeptic's Guide to Objective Chance." *Mind* 116:549–596.

Howson, Colin. 2000. *Hume's Problem: Induction and the Justification of Belief*. Oxford University Press.

[63] It's worth noting that an impermissive Bayesian could reject the claim that her fortunes are tied to those of Indifference Principles, and instead hold that there are a number of heterogeneous constraints on rational priors, the combination of which leads to a unique permissible prior. This position would avoid many of the problems surrounding Indifference Principles.

[64] For helpful comments and discussion, I'd like to thank Phil Bricker, Maya Eddon, Darin Harootunian, Sophie Horowitz, Tom Kelly, Miriam Schoenfield, Katia Vavova, Jonathan Vogel, Jonathan Weisberg, Roger White, the Northwest Time and Rationality conference, the MIT epistemology reading group, and two anonymous referees.

Howson, Colin and Peter Urbach. 2005. *Scientific Reasoning: The Bayesian Approach*. 3rd ed. Open Court Publishing Company.

Jaynes, Edwin. 1983. *Papers on Probability, Statistics and Statistical Physics*. Dordrecht: Reidel.

Kelly, Thomas. 2012. How to be a Permissivist. In *APA Symposium*. Blackwell.

Lewis, David. 1986. A Subjectivist's Guide to Objective Chance. In *Philosophical Papers, Vol. 2*. Oxford University Press.

Lewis, David. 1994. "Humean Supervenience Debugged." *Mind* 103:473–490.

Meacham, Christopher J G. 2005. "Three Proposals Regarding a Theory of Chance." *Philosophical Perspectives* 19:281–307.

North, Jill. 2010. "An Empirical Approach to Symmetry and Probability." *Studies in History and Philosophy of Modern Physics* 41:27–40.

Schaffer, Jonathan. 2007. "Deterministic Chance?" *The British Journal for the Philosophy of Science* 58:113–140.

Sider, Theodore. 2011. "Writing the Book of the World." Forthcoming from Oxford University Press.

van Fraassen, Bas. 1984. "Belief and the Will." *Journal of Philosophy* 81:235–256.

van Fraassen, Bas. 1989. *Laws and Symmetry*. Oxford University Press.

Weisberg, Jonathan. 2011. Varieties of Bayesianism. In *Handbook of the History of Logic*, ed. Dov Gabbay, Stephen Hartmann and John Woods. Vol. 10 North Holland pp. 477–552.

White, Roger. 2005. "Epistemic Permissiveness." *Philosophical Perspectives* 19:445–459.

White, Roger. 2010. Evidential Symmetry and Mushy Credences. In *Oxford Studies in Epistemology*, ed. Tamar Szabó Gendler and John Hawthorne. Vol. 3 Oxford University Press pp. 161–186.

Zimmerman, Michael J. 2006. "Is Moral Obligation Objective or Subjective?" *Utilitas* 18:329–361.