

Evidential Incomparability and the Principle of Indifference

Martin Smith

The *Principle of Indifference* was once regarded as a linchpin of probabilistic reasoning, but has now fallen into disrepute as a result of the so-called *problem of multiple of partitions*. In ‘Evidential symmetry and mushy credence’ Roger White suggests that we have been too quick to jettison this principle and argues that the problem of multiple partitions rests on a mistake. In this paper I will criticise White’s attempt to revive the Principle of Indifference. In so doing, I will argue that what underlies the problem of multiple partitions is a fundamental tension between the Principle of Indifference and the very idea of *evidential incomparability*.

I. MULTIPLE PARTITIONS

Say that two propositions P and Q are *evidentially symmetric* ($P \approx Q$) iff one’s evidence does not support P more strongly than Q ($\sim (P \gg Q)$) and does not support Q more strongly than P ($\sim (P \ll Q)$). The Principle of Indifference (POI) is a way of capturing the idea that I should not discriminate, in my credence assignments, between propositions that are evidentially symmetric. Suppose $\{P_1 \dots P_n\}$ is a partition – that is, a set of propositions that, given my evidence, are mutually exclusive and jointly exhaustive – and suppose that Cr is a rational credence assignment – a way that I can rationally assign credence to propositions. POI states that, if $P_1 \approx P_2 \approx \dots \approx P_n$ then, for all i , $1 \leq i \leq n$ $Cr(P_i) = 1/n$. According to POI, if I’m confronted by a set of n propositions that, given my evidence, are mutually exclusive and jointly exhaustive and my evidence provides no more support for any one of these over any other, then the only rational thing to do is assign a credence of $1/n$ to each. Though it was once widely accepted, POI has long been thought by epistemologists to fall prey to the *problem of multiple partitions*.

Suppose I find myself in a factory that manufactures square plates¹. I know that these squares must have a side length of less than 2 feet – but this is all the information that I have about them. Suppose a new square is just about to roll off the production line. What should my credence be in the proposition that the square has a side length of less than 1 foot? There are two possibilities here – either the square has a side length of less than 1 foot or the square has a side length between 1 and 2 feet – and my evidence no more supports one than the other. According to POI, if I am rational then my credence in the proposition that the square has a side length of less than 1 foot must be $1/2$.

¹ This example is a variant on van Fraassen’s ‘cube factory’ (van Fraassen, 1989). Cases of this general sort were first described by Bertrand (1889).

If the squares produced by the factory all have side lengths of less than 2 feet, it follows, of course, that they all have areas of less than 4 square feet. What should my credence be in the proposition that the next square to roll off the production line has an area of less than 1 square foot? Here there are four possibilities – either the square has an area of less than 1 square foot or it has an area between 1 and 2 square feet or it has an area between 2 and 3 square feet or it has an area between 3 and 4 square feet – and my evidence no more supports any one of these than it does any other. According to POI, if I am rational then my credence in the proposition that the square has an area of less than 1 square foot must be $1/4$. But the proposition that the square has an area of less than 1 square foot is, of course, *equivalent* to the proposition that the square has a side length of less than 1 foot. These are just two ways of describing the very same condition. POI has generated two conflicting pieces of advice.

What we have, in effect, are two equivalent propositions embedded in different partitions:

$$L_1: 0 \leq \text{length} < 1 \text{ ft.} \quad \leftrightarrow \quad A_1: 0 \leq \text{area} < 1 \text{ sq. ft.}$$

$$L_2: 1 \leq \text{length} < 2 \text{ ft.} \quad A_2: 1 \leq \text{area} < 2 \text{ sq. ft.}$$

$$A_3: 2 \leq \text{area} < 3 \text{ sq. ft.}$$

$$A_4: 3 \leq \text{area} < 4 \text{ sq. ft.}$$

Given that $L_1 \approx L_2$ and that $A_1 \approx A_2 \approx A_3 \approx A_4$ we can quickly derive an absurdity using POI. Call this the *multiple partitions argument*:

(i)	$L_1 \approx L_2$	Premise
(ii)	$A_1 \approx A_2 \approx A_3 \approx A_4$	Premise
(iii)	$\text{Cr}(L_1) = 1/2$	i, POI
(iv)	$\text{Cr}(A_1) = 1/4$	ii, POI
(v)	$\text{Cr}(L_1) = \text{Cr}(A_1)$	$L_1 \leftrightarrow A_1$
(vi)	$1/2 = 1/4$	iii, iv, v

Despite occasional claims to the contrary, the multiple partitions argument doesn't show that POI is *internally* contradictory or inconsistent. Premises (i) and (ii) are pivotal in deriving the absurd (vi) and, though they do appear very plausible, they are hardly logical truths. In 'Evidential symmetry and mushy credence' (2010) Roger White argues that these two premises are less innocent than they first appear – especially when taken in combination. As White shows, it is possible to derive an absurdity from (i) and (ii) *alone*, provided we make use of two further principles that he describes as 'obviously true':

Equivalence: If P and Q are known to be equivalent then $P \approx Q$.

Transitivity: If $P \approx Q$ and $Q \approx R$ then $P \approx R$.

With these principles in place the derivation – which we might call *White's argument* – can proceed as follows:

(i)	$A_2 \approx A_1$	Premise
(ii)	$A_1 \approx L_1$	Equivalence
(iii)	$A_2 \approx L_1$	i, ii, Transitivity
(iv)	$L_1 \approx L_2$	Premise
(v)	$A_2 \approx L_2$	iii, iv, Transitivity
(vi)	$L_2 \approx (A_2 \vee A_3 \vee A_4)$	Equivalence
(vii)	$A_2 \approx (A_2 \vee A_3 \vee A_4)$	v, vi, Transitivity

The conclusion does indeed seem absurd – given that my evidence leaves A_3 and A_4 open, it must provide more support for $A_2 \vee A_3 \vee A_4$ than for A_2 alone. That is, my evidence must give me more reason to think that the square has an area between 1 and 4 square feet than that the square has an area between 1 and 2 square feet. But if (i) and (ii) lead to absurdity when combined with obviously true principles, then the fact that they lead to absurdity when combined with POI should hardly be thought to cast doubt upon POI.

It is crucial to White's overall argumentative strategy that Equivalence and Transitivity are genuinely uncontroversial principles. While Equivalence may be relatively safe, I will argue here that Transitivity is in fact highly suspect, both in general and in application to the particular case under consideration². As such, White has not succeeded in showing that the combination of (i) and (ii) is in any way absurd and has not succeeded in disarming the original multiple partitions argument against POI. My first aim here is simply to reinstate this argument. My second aim is to look in more detail at the relationship between POI and Transitivity. As it turns out, the viability of both of these principles hinges upon a deeper issue – namely, whether there is such a thing as *evidential incomparability*. If there is such a thing, then both POI and Transitivity must be abandoned. Before making good on these claims, a little more stage setting is required.

II. MUSHY CREDENCE

It is not difficult to appreciate why POI was once so highly regarded. It is very natural to reason in something like the following way: Rational credences should reflect one's evidence. But, if I'm confronted by a partition of n propositions that are all evidentially symmetric, then assigning a credence of $1/n$ to each would seem to be the only option that

² Novack (2010) objects to the transitivity of evidential symmetry by appealing to the apparent failure of the transitivity of indistinguishability. Suppose I'm shown a sequence of colour samples ranging from, say, red to blue, such that adjacent samples are pairwise indistinguishable and one of the samples matches the colour of a getaway car that I saw speeding away from a robbery. If I'm shown any two adjacent samples, my evidence, plausibly, no more supports the proposition that the first matches the colour of the car than the proposition that the second matches the colour of the car, since I can't tell them apart. But if we consider the proposition that the first (red) colour in the sequence matches the colour of the car and the proposition that the last (blue) colour in the sequence matches the colour of the car, my evidence could well support one of these more strongly than the other. Novack attributes this example to Branden Fitelson. White also briefly considers this kind of objection to Transitivity and attributes it both to Fitelson and to Elliot Sober. Whatever we make of this example, though, it seems to be of limited value in responding to White's argument. After all, White does not require that Transitivity hold universally – merely that it hold in the square factory and in other cases like it. And the present considerations don't give us any obvious reason to question that.

could possibly do so. The alternatives, after all, will involve assigning a *higher* credence to some propositions and a *lower* credence to others – but, by stipulation, I don't have any evidential basis for doing this. This reasoning is, I think, most compelling against the background of the orthodox Bayesian assumption that a rational credence assignment must meet the conditions for a *probability* function. If that's right and I'm rational, then the credences I assign to the n propositions in the partition must be precise numerical values and must sum to 1. Given these constraints, assigning $1/n$ credence to each proposition is indeed the only possible assignment that doesn't involve favouring some propositions over others.

Whatever the merits of this orthodox Bayesian view of credences, a number of philosophers have recently become attracted to a different sort of picture – a picture on which, as well as taking precise values, rational credences can potentially be 'spread out' over a range of different values. As it's variously put, credences can be 'imprecise' or 'thick' or 'mushy' (see, for instance, Joyce, 2005, Weatherson, 2007, Sturgeon, 2008). Given this sort of picture, the above reasoning seems rather less persuasive. If I'm confronted by a partition of n propositions that are all evidentially symmetric then, rather than assigning each proposition a precise credence of $1/n$, I could make my credence in each proposition mushy instead – that is, make my credence in each proposition cover some appropriate spread of values. And there are various ways of assigning mushy credences to each proposition without favouring some over others.

The idea that rational credences should reflect evidence can, perhaps, be distilled into the following principle, which we might call the *evidential constraint* (EC): $\text{Cr}(P) > \text{Cr}(Q)$ iff $P \gg Q$. According to EC, it is rational for one to assign a higher credence to P than to Q iff one's evidence supports P more strongly than it supports Q . If we assume that rational credences must be precise numerical values and must sum to 1 across a partition, then EC entails POI. But if we allow that credences can be mushy, then a gap opens between EC and POI. There are ways of assigning mushy credences to the members of an evidentially symmetric partition that are consistent with the recommendations of EC. But no such assignment is consistent with the recommendations of POI.

Different philosophers, I suspect, understand the notion of mushy credence in rather different ways. Most, though, make use of a certain formal model on which a rational credence assignment is captured not by a *single* probability function but by a *set* of probability functions Γ , often termed a *representor*. On this model, the credence assigned to a particular proposition will be equal to the set of values assigned to that proposition by the probability functions in one's representor – $\text{Cr}(P) = \{i \mid \exists \text{Pr}_x \in \Gamma, \text{Pr}_x(P) = i\}$. The composition of a representor is often constrained in such a way as to ensure that this set of values forms a real interval – though such a constraint won't assume much significance for present purposes³. The credence assigned to a proposition will be precise in case all of the functions in one's representor assign it the same probability value and it will be mushy

³ A representor is often required to be *convex* – that is, closed under the operation of taking weighted averages of probability functions. More formally, a representor Γ is convex just in case for every $\text{Pr}_x \in \Gamma$ and $\text{Pr}_y \in \Gamma$, $(\alpha\text{Pr}_x + (1 - \alpha)\text{Pr}_y) \in \Gamma$, for all real numbers α , $0 \leq \alpha \leq 1$. The weighted average of two probability functions is a function that assigns to each proposition the weighted average of the two values assigned by the two functions. As can be clearly seen, the values assigned to a proposition by the functions in a convex representor must themselves be closed under the taking of weighted averages and, thus, form a real interval.

otherwise. One's credence in a proposition will be maximally mushy just in case, for every real number in the unit interval, there is a function in one's representor that assigns that number to the proposition. This might be thought to represent a state of total agnosticism or suspension of judgment about the proposition⁴.

Given this model, one option that is open to us when assigning credences to the members of a partition $\{P_1...P_n\}$ is to be totally agnostic about each. Suppose that, for every string of n real numbers that sum to 1, there is some function in one's representor that assigns these numbers, respectively, to the members of the partition. On the present model, what it is for a proposition P to be assigned a higher credence than a proposition Q is for some function in one's representor to assign P a higher probability than Q and for no function in one's representor to assign P a lower probability than Q – that is, for some function in one's representor to reckon P more probable than Q and for every function in one's representor to reckon P at least as probable as Q :

$$\text{Cr}(P) > \text{Cr}(Q) \text{ iff } (\exists \text{Pr}_x \in \Gamma, \text{Pr}_x(P) > \text{Pr}_x(Q)) \wedge (\forall \text{Pr}_x \in \Gamma, \text{Pr}_x(P) \geq \text{Pr}_x(Q)).$$

Assigning a total agnostic credence to each member of $\{P_1...P_n\}$ would not, then, involve assigning a higher credence to some members and a lower credence to others. If the members of $\{P_1...P_n\}$ are evidentially symmetric, then the supposition that such an assignment is rational is consistent with EC, though it violates POI which demands that my credence in each member be precisely $1/n$. This, of course, is just one possible mushy credence assignment that has these properties.

If we combine the above analysis of credal comparisons with EC then we derive the following analyses of *evidential* comparisons:

$$P \gg Q \text{ iff } (\exists \text{Pr}_x \in \Gamma, \text{Pr}_x(P) > \text{Pr}_x(Q)) \wedge (\forall \text{Pr}_x \in \Gamma, \text{Pr}_x(P) \geq \text{Pr}_x(Q))$$

$$P \ll Q \text{ iff } (\exists \text{Pr}_x \in \Gamma, \text{Pr}_x(P) < \text{Pr}_x(Q)) \wedge (\forall \text{Pr}_x \in \Gamma, \text{Pr}_x(P) \leq \text{Pr}_x(Q))$$

From these we can, in turn, derive an analysis of evidential symmetry. What it is for P and Q to be evidentially symmetric, recall, is for one's evidence to neither support P more strongly than Q or Q more strongly than P – that is, for neither of the above conditions to be met:

⁴ Orthodox Bayesians famously demand that the only way to rationally update one's credence function is by *conditionalising* on new evidence – if one receives evidence E and no further evidence and Cr and Cr_E are one's credence functions before and after the receipt of this evidence then, provided $\text{Cr}(E) > 0$ and one is rational, $\text{Cr}_E(P) = \text{Cr}(P | E) = \text{Cr}(P \wedge E) / \text{Cr}(E)$ for any P . This rule has an obvious analogue within the mushy credence model that we might adopt – the only way to rationally update one's representor is by conditionalising *every* function therein on new evidence. If, however, one is in a state of total agnosticism about a proposition P then, given certain assumptions, conditionalising on a proposition E that is assigned a positive value by every function in the representor will leave the distribution of values assigned to P untouched – that is, it will leave one in a state of total agnosticism about P . If one is in a state of total agnosticism about P then there is a function in one's representor that assigns P a probability of 0. This function will continue to assign P a probability of 0 even after being conditionalised on E . If one is in a state of total agnosticism about P then there is a function in one's representor that assigns P a probability of 1. This function will continue to assign P a probability of 1 even after being conditionalised on E . The convexity constraint will then ensure that the functions in one's representor continue to assign all real values in the unit interval to P . This would appear to make agnosticism into a very unappealing prospect – a state that, once entered into, may never be rationally escaped. For this reason, those who adopt the mushy credence model – particularly those who adopt it as a way of doing justice to the attitude of agnosticism – often use alternative, more permissive, update rules (see for instance Weatherston, 2007).

$$\begin{aligned}
P \approx Q \quad \text{iff} \quad & ((\forall \text{Pr}_x \in \Gamma, \text{Pr}_x(P) \leq \text{Pr}_x(Q)) \vee (\exists \text{Pr}_x \in \Gamma, \text{Pr}_x(P) < \text{Pr}_x(Q))) \wedge \\
& ((\forall \text{Pr}_x \in \Gamma, \text{Pr}_x(P) \geq \text{Pr}_x(Q)) \vee (\exists \text{Pr}_x \in \Gamma, \text{Pr}_x(P) > \text{Pr}_x(Q))) \\
\text{iff} \quad & ((\forall \text{Pr}_x \in \Gamma, \text{Pr}_x(P) \leq \text{Pr}_x(Q)) \wedge (\forall \text{Pr}_x \in \Gamma, \text{Pr}_x(P) \geq \text{Pr}_x(Q)) \vee \\
& ((\exists \text{Pr}_x \in \Gamma, \text{Pr}_x(P) < \text{Pr}_x(Q)) \wedge (\exists \text{Pr}_x \in \Gamma, \text{Pr}_x(P) > \text{Pr}_x(Q))) \\
\text{iff} \quad & ((\forall \text{Pr}_x \in \Gamma, \text{Pr}_x(P) = \text{Pr}_x(Q)) \vee (\exists \text{Pr}_x, \text{Pr}_y \in \Gamma, \text{Pr}_x(P) < \text{Pr}_x(Q) \wedge \text{Pr}_y(P) \\
& > \text{Pr}_y(Q))
\end{aligned}$$

Less formally, P and Q are evidentially symmetric iff *either* every function in one's representor reckons P and Q to be equally probable – assigns them the same probability value – or the functions in one's representor disagree as to whether P is more probable than Q or Q is more probable than P – some functions in one's representor reckon P to be more probable than Q while others reckon Q to be more probable than P. This, then, is how we ought to conceive of evidential symmetry on the present model, given EC.

As can be easily demonstrated, though, this is *not* a transitive relation. Suppose one's representor consists of just two probability functions Pr_1 and Pr_2 . Suppose that:

$$\begin{array}{ll}
\text{Pr}_1(P) = 3/10 & \text{Pr}_2(R) = 3/10 \\
\text{Pr}_1(Q) = 2/10 & \text{Pr}_2(P) = 2/10 \\
\text{Pr}_1(R) = 1/10 & \text{Pr}_2(Q) = 1/10
\end{array}$$

According to Pr_1 P has a higher probability than R and according to Pr_2 R has a higher probability than P. Thus, P and R are evidentially symmetric. According to Pr_2 R has a higher probability than Q and according to Pr_1 Q has a higher probability than R. Thus, R and Q are evidentially symmetric. But, according to *both* Pr_1 and Pr_2 , P has a higher probability than Q. Thus, P and Q are not evidentially symmetric – on the contrary, one's evidence supports P more strongly than Q. We have it that $P \approx R$, $R \approx Q$ and $P \gg Q$ all at once⁵.

If we adopt the present model of mushy credence, and combine it with EC, then Transitivity fails. Not only this, but on one natural way of implementing the model, Transitivity will fail precisely in multiple partitions cases, such as the square factory case described above. Given my evidence in the case, other than taking the view that the square has a side length of less than 2 feet (and an area of less than 4 square feet) it looks as though I should, rationally, be in a state of total agnosticism about its length and area. That is, for each of L_1, L_2, A_1, A_2, A_3 and A_4 and for each real number n , $0 \leq n \leq 1$, there will be some function in my representor that assigns n to the proposition. But for every function in my representor, of course, the probability values assigned to L_1 and L_2 will sum to 1 as will the probability values assigned to A_1, A_2, A_3 and A_4 . Furthermore, every function in my representor will assign the same probability to the equivalent propositions L_1 and A_1 and to the equivalent propositions L_2 and $(A_2 \vee A_3 \vee A_4)$. As the square factory case makes vivid, no *single* probability function can be indifferent with respect to both side length and area – a

⁵ If one wishes to impose a convexity constraint upon representors, then we could simply consider the closure of $\{\text{Pr}_1, \text{Pr}_2\}$ under the taking of weighted averages. As can be easily checked, moving to this representor will alter none of the above verdicts.

function that is indifferent with respect to side length will be biased with respect to area and a function that is indifferent with respect to area will be biased with respect to side length. But a *representor* can, after a fashion, be indifferent with respect to both side length and area – namely, by including functions of both sorts. The representor described will include functions of both sorts, and a good many besides⁶.

Given this representor, we have it that $L_1 \approx L_2$, since there will be functions that assign L_1 a higher probability than L_2 and functions that assign L_2 a higher probability than L_1 . We have it that that $A_1 \approx A_2 \approx A_3 \approx A_4$ since, for any one of these propositions, there will be functions that assign it a higher probability than the other three and functions that assign it a lower probability than the other three. We also have it that that $L_1 \approx A_1$ since every function will assign these two propositions the same probability and that $L_2 \approx (A_2 \vee A_3 \vee A_4)$ since every function will assign these two propositions the same probability. Returning to White's argument, as formulated at the end of the last section, the first application of Transitivity will be sound – it will turn out that $A_2 \approx L_1$ – but the next application of Transitivity will fail – it will *not* turn out that $A_2 \approx L_2$ or that $A_2 \approx (A_2 \vee A_3 \vee A_4)$. On the contrary, we have it that $A_2 \ll L_2$ and $A_2 \ll (A_2 \vee A_3 \vee A_4)$ since some functions will assign L_2 and $(A_2 \vee A_3 \vee A_4)$ a higher probability than A_2 and none will assign A_2 a higher probability than L_2 and $(A_2 \vee A_3 \vee A_4)$.

III. EQUAL VS INCOMPARABLE SUPPORT

For a defender of the mushy credence model, Transitivity is false and premises (i) and (ii) in the multiple partitions argument are in no tension with one another. For a defender of the mushy credence model, the multiple partitions argument against POI stands. But this, in and of itself, may be of little consequence. After all, this is just a model – and needn't be treated as sacrosanct. We could easily build another sort of model that makes quite different predictions. In the later sections of 'Evidential symmetry and mushy credence' White offers a number of criticisms of the mushy credence model and, if these criticisms are good, then perhaps we needn't take its predictions all that seriously. Ultimately, though, I think it matters little what we make of this model or indeed of the idea of mushy credence in general. For the model merely serves to *illustrate* something that is, I think, very plausible in its own right – namely, there are two distinct ways in which a pair of propositions might qualify as evidentially symmetric. First, the propositions might be *equally strongly supported* by one's

⁶ What if my representor included some functions that were indifferent with respect to side length and some functions that were indifferent with respect to area and *no* functions of any other kind? In this case I wouldn't count as being totally agnostic about the side length and area of the plate. My credence in L_1 for instance would be {0.25, 0.5} while my credence in L_2 would be {0.5, 0.75}. Could I rationally adopt such a credal state in the square factory case? One concern about this representor is that would violate the convexity constraint mentioned in footnote 3. But even if we put aside such worries, it's not clear that such a representor would be a legitimate response to the evidence I have available. After all, I have no reason to think that the squares produced by this factory must be evenly distributed with respect to side length or with respect to area or, indeed, with respect to any other parameter that might strike me as being natural. As far as my evidence is concerned, the squares produced by the factory could have any distribution whatever, provided none of them have side lengths in excess of two feet. I won't explore this further here.

evidence (this is what's captured by the first disjunct in the above analysis of evidential symmetry) or, second, the propositions might be *incomparable with respect to evidential support* (this is what's captured by the second disjunct in the above analysis of evidential symmetry). Whatever one makes of the particulars of the mushy credence model, this prediction is a sound one.

White himself provides two contrasting examples of evidential symmetry – one that is naturally understood as a case of equal evidential support and one that is naturally understood as a case of evidential incomparability. Suppose I'm about to draw a marble from an urn. Suppose, in the first case, that I know that the urn contains only five black marbles and five white marbles and has just been shaken up. Suppose, in the second case, that I know that the urn contains some marbles but I have no further information about its contents. In the first case, it seems natural to say that my evidence provides equally strong support for the proposition that I'll draw a black marble and the proposition that I'll draw a white marble. We can, perhaps, persuade ourselves to say something similar about the second case as well – but we are left with an awkward sense that we're papering over some kind of significant difference between the two cases. In the first case, my evidence provides a certain quantifiable amount of support for the proposition that I'll draw a white marble and a certain quantifiable amount of support for the proposition that I'll draw a black marble and these amounts happen to be equal. In the second case, my evidence provides *some* support for the proposition that I'll draw a white marble and for the proposition that I'll draw a black marble, at least in so far as it's consistent with each. But any attempt to quantify this support seems quite out of place.

Here, perhaps, is one way to reinforce these impressions. White marbles are not all identical – they come, for instance, in subtly different shades of white. Let white-7 be one of these shades. Presumably, my evidence provides stronger support for the proposition that I'll draw a white marble than the proposition that I'll draw a white-7 marble. As such, if my evidence really does provide equally strong support for the proposition that I'll draw a black marble and the proposition that I'll draw a white marble then it must provide *stronger* support for the proposition that I'll draw a black marble than the proposition that I'll draw a white-7 marble. In the first case, this would indeed seem to be so. In the first case, I know that there are at least as many black marbles in the urn as there are white-7 marbles and there could well be more. Thus, my evidence does give me more reason to think that I'll draw a black marble than a white-7 marble. But in the second case my evidence doesn't support the proposition that I'll draw a black marble more strongly than the proposition that I'll draw a white-7 marble. How could it? My evidence, in the second case, gives me absolutely no reason to think that the urn contains more black marbles than white-7 marbles – it's completely silent on all such matters. Thus, the second marble case is not really a case of equal evidential support⁷.

⁷This argument doesn't force one to admit the existence of evidential incomparability. One might, for instance, insist that the proposition that I'll draw a black marble *is* more strongly supported than the proposition that I'll draw a white-7 marble, on the grounds that white-7 is a more precise, circumscribed colour than black – almost as though there's some objective 'colour space' in which black occupies a larger region than white-7. This is not an incoherent line of thought – but it does strike me as difficult to maintain. Even if we can make sense of an objective colour space, the fact remains that I have no reason at all to think that the distribution of marbles in the

The square factory is, of course, closely analogous to the second marble case above. Both are cases in which I'm speculating about the properties of an item, randomly selected from a population, without any information about the frequency with which the relevant properties occur in the population. In the square factory, I have no basis for comparing the proposition that the square has a side length between 0 and 1 foot and the proposition that the square has a side length between 1 and 2 feet and, similarly, no basis for comparing the propositions that the square has an area between 0 and 1 square foot, between 1 and 2 square feet, between 2 and 3 square feet and between 3 and 4 square feet. As noted, I have no information about how these properties are distributed in the underlying population.

Now suppose that, unlike in the original case, I do have some information of this kind. Suppose I know that the factory is set up in such a way that 50% of the squares it produces will have a side length of between 0 and 1 foot and 50% will have a side length of between 1 and 2 feet. In this case the proposition that the square has a side length between 0 and 1 foot and the proposition that the square has a side length between 1 and 2 feet would be equally supported by my evidence. But the proposition that the square has an area of between 0 and 1 square foot would no longer be evidentially symmetric with, say, the proposition that the square has an area of between 1 and 2 square feet – rather, it would be more strongly supported. After all, I would know that the factory will produce at least as many squares of the first type as the second type and could well produce more. In this revised case, premise (i) of the multiple partitions argument would still be true but premise (ii) would be false.

Suppose, on the other hand, that I know the factory to be set up in such a way that 25% of the squares it produces will have an area between 0 and 1 square foot, 25% will have an area between 1 and 2 square feet, 25% will have an area between 2 and 3 square feet and 25% will have an area between 3 and 4 square feet. In this revised case, the propositions that the square has an area between 0 and 1 square foot, that the square has an area between 1 and 2 square feet, that the square has an area between 2 and 3 square feet and that the square has an area between 3 and 4 square feet would all, plausibly, be equally supported. But the proposition that the square has a side length between 0 and 1 foot would no longer be evidentially symmetric with the proposition that the square has a side length between 1 and 2 feet – rather, it would be less strongly supported. After all, I would know that the factory will produce more squares of the second type than the first. While premise (ii) of the multiple partitions argument would be true, premise (i) would be false.

In the square factory, the propositions in each partition are described as being evidentially symmetric – and the description seems to fit. But these evidential symmetry relations must rest on underlying relations of evidential incomparability, rather than relations of equal evidential support. If we adjust the case so that the propositions in one partition are equally supported, then the evidential symmetries in the other partition are broken. We need the notion of evidential incomparability in order to make sense of the original description of the square factory case.

The idea that evidential symmetry can arise either through equal evidential support or through evidential incomparability is not, then, a mere artefact of the mushy credence model.

urn should, in any way, conform to it and, thus, no reason at all to think that it has any bearing on the question of what colour marble I'll draw.

Arguably, it is an idea that *any* adequate model of evidential support ought to capture. But, once we concede that evidential symmetry is the disjunction of equal evidential support and evidential incomparability, then, as far as Transitivity is concerned, the game is effectively up. While equal evidential support would appear to be a transitive relation⁸, evidential incomparability is clearly not. Incomparability relations are not, in general, transitive. Suppose one thinks, for instance, that abstract expressionist paintings and nineteenth century landscape paintings are incomparable with respect to beauty. One would then think that Constable's *Hay Wain* cannot be compared with Pollock's *Blue Poles* which, in turn, cannot be compared with Constable's *Flatford Mill*. But one is hardly obliged to conclude that *Hay Wain* and *Flatford Mill* cannot be compared.

What goes for aesthetic incomparability, also goes for evidential incomparability. In the second marble case, the proposition that I'll draw a white marble cannot be compared with the proposition that I'll draw a black marble. I simply don't know how the number of white marbles in the urn compares with the number of black marbles. Further, the proposition that I'll draw a black marble cannot be compared with the proposition that I'll draw a white-7 marble. I simply don't know how the number of black marbles in the urn compares with the number of white-7 marbles. But we're hardly obliged to conclude that the proposition that I'll draw a white marble and the proposition that I'll draw a white-7 marble cannot be compared. One thing that I *do* know about the contents of the urn is that it must contain at least as many white marbles as white-7 marbles and could well contain more. This is true irrespective of what the urn contains.

In the square factory, the proposition that the square has an area between 1 and 2 square feet cannot be compared with the proposition that the square has a side length between 0 and 1 foot. I simply don't know how the number of produced squares with an area of between 1 and 2 square feet will compare with the number of produced squares with a side length between 0 and 1 foot. Further, the proposition that the square has a side length between 0 and 1 foot cannot be compared with the proposition that the square has a side length between 1 and 2 feet. I simply don't know how the number of produced squares with a side length between 0 and 1 foot will compare with the number of produced squares with a side length between 1 and 2 feet. But we're hardly obliged to conclude that the proposition that the square has an area between 1 and 2 square feet and the proposition that the square has a side length between 1 and 2 feet cannot be compared. One thing that I do know is that the factory must produce at least as many squares with a side length between 1 and 2 feet as squares with an area between 1 and 2 square feet, and could well produce more. This is true irrespective of how the factory is set up.

The claim that the evidential incomparability relation is not transitive can, in fact, be given something like a formal proof – it follows just from the suppositions that it is

⁸ If counterexamples to the transitivity of indistinguishability do serve as counterexamples to the transitivity of evidential symmetry, then they will presumably also serve as counterexamples to the transitivity of equal evidential support. If the proposition that the first of two adjacent colour samples matches the colour of the car is evidentially symmetric with the proposition that the second of two adjacent colour samples matches the colour of the car, then this would have to be because they are equally strongly supported – they don't look to be evidentially incomparable. The transitivity of equal evidential support could perhaps be preserved, in the face of such cases, if we are prepared to give up on the idea that discrepancies in evidential support, no matter how small, must always be distinguishable. I won't pursue this further here.

symmetric, irreflexive and instantiated. If P is evidentially incomparable with Q then Q must be evidentially incomparable with P. If evidential incomparability were transitive then P would have to be evidentially incomparable with *itself* – but this, presumably, is impossible. Such cases won't, however, provide counterexamples to the transitivity of evidential symmetry. Although the transitivity of evidential incomparability will fail in cases with this structure, the transitivity of evidential symmetry will be rescued by its *other disjunct* – presumably P will always be just as strongly supported as itself. To refute the transitivity of evidential symmetry, we need cases in which one's evidence more strongly supports the proposition at one end of a chain of evidential incomparability relations than the proposition at the other. But anyone who accepts the possibility of evidential incomparability should accept the existence of such cases – the second marble case and the square factory case are of exactly this kind.

Transitivity is only acceptable on the assumption that there is no such thing as evidential incomparability and that the only way in which propositions can be evidentially symmetric is by enjoying equally strong evidential support. The multiple partitions argument, on the other hand, requires the assumption that there *is* such a thing as evidential incomparability – otherwise (i) and (ii) cannot be jointly accepted. It should come as no particular surprise, then, that attempting to combine Transitivity with (i) and (ii) should lead to absurdity. None of these observations, though, brings any comfort to a defender of POI. The final dialectical state of play is just this: If we reject the possibility of evidential incomparability, then POI and Transitivity remain viable – the premises of the multiple partitions argument and of White's argument will be inconsistent. But if we accept the possibility of evidential incomparability, then both arguments engage – White's argument serves as a powerful reductio of Transitivity and the multiple partitions argument serves as a powerful reductio of POI, just as epistemologists have thought all along.

References

- Bertrand, J. (1889) *Calul des Probabilitiés* (Paris: Gaultier Villers)
- Joyce, J. (2005) 'How probabilities reflect evidence' *Philosophical Perspectives* v19 pp153-178
- Novack, G. (2010) 'A defense of the principle of indifference' *Journal of Philosophical Logic* v39, pp655-678
- Sturgeon, S. (2008) 'Reason and the grain of belief' *Noûs* v42, pp139-165
- Van Fraassen, B. (1989) *Laws and Symmetry* (Oxford: Oxford University Press)
- Weatherson, B. (2007) 'The Bayesian and the dogmatist' *Proceedings of the Aristotelian Society* v107, pp169-185
- White, R. (2010) 'Evidential symmetry and mushy credence' in Szabo Gendler, T. and Hawthorne, J. eds. *Oxford Studies in Epistemology Volume 3* (New York: Oxford University Press)