

OXFORD STUDIES IN EPISTEMOLOGY

Volume 3

Edited by

Tamar Szabó Gendler and John Hawthorne

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

© The several Contributors 2010

The moral rights of the authors have been asserted
Database right Oxford University Press (maker)

First published 2010

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose the same condition on any acquirer

British Library Cataloguing in Publication Data

Data available

Library of Congress Cataloging in Publication Data

Data available

Typeset by SPI Publisher Services, Pondicherry, India

Printed in Great Britain
on acid-free paper by
MPG Books Group,
Bodmin and King's Lynn

ISBN 978-0-19-958408-6 (hbk.)

ISBN 978-0-19-958409-3 (pbk.)

1 3 5 7 9 10 8 6 4 2

7. Evidential Symmetry and Mushy Credence

Roger White

Often our evidence supports some hypotheses over others. To the extent that it does we can put more confidence in those that enjoy greater evidential support. But sometimes it doesn't. Sometimes all that we have to go on gives us no guidance as to which of several alternatives is true (perhaps because we have little if anything to go on at all). In such cases the Principle of Indifference (a.k.a. Principle of Insufficient Reason) tells us to match our confidence to the symmetry of our evidential situation. If our confidence is best modeled by a standard probability function this means that we are to distribute our subjective probability or credence sharply and evenly over possibilities among which our evidence does not discriminate.

Once thought to be *the* central principle of probabilistic reasoning by great thinkers like Laplace (1825/1951), the Principle of Indifference has fallen on hard times. These days it is commonly dismissed as an old-fashioned item of confusion.¹ While I certainly haven't found my way through all the difficulties in this area, I want to suggest that we need to rethink the matter. I will argue that the objections to the principle are not as devastating as they appear, and that there is a compelling case in support of the basic idea, which cannot be ignored. Lastly, I'll consider the most important alternative to the principle, which involves the idea that one's credence should not always be 'sharp'.

1. EVIDENTIAL SYMMETRY

Let's say that

Propositions p and q are **evidentially symmetrical** (I'll write this as $p \approx q$) for a subject if his evidence no more supports one than the other.

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

Thanks to audiences at University of Vermont, the University of North Carolina at Chapel Hill, Princeton, Australian National University, Massachusetts Institute of Technology, and the Formal Epistemology Workshop at the University of Wisconsin at Madison for feedback. In particular I've benefited greatly from conversation or correspondence with Rachael Briggs, David Christensen, Adam Elga, David Edin, Branden Fitelson, Alan Hajék, Caspar Hare, Matt Kotzen, Sarah Moss, James Overton, Agustin Rayo, Susanna Rinard, Teddy Seidenfeld, Mike Smithson, Scott Sturgeon, Elliott Sober, and Bas van Fraassen.

¹ Recent critics include van Fraassen (1989), Joyce (1995), Strevens (1998), Gillies (2000), North (MS), and Sober (2003).

‘evidence’, if there is such. We might say: p and q are evidentially symmetrical for you iff you have *no more reason* to suppose that p is true than that q is, or vice versa.

1.1. Sources of Evidential Symmetry

Let’s distinguish different ways that evidential symmetry can arise. I’m wondering if the marble taken from this urn is black or white. I might have a rich body of data relevant to each proposition that bears on them in the same way:

Case 1: I know that the marble has been selected from a shaken urn containing just five white and five black balls.

Or I might have no relevant evidence either way:

Case 2: I have no idea how or why this marble was chosen, or the constitution of the urn from which it came (for all I know it might contain unequal numbers of white and black balls, but I have no evidence concerning which color it has more of, if any).

In each case I have no more reason to suppose that the ball is black than that it is white. In the first we might say that I know that the *objective chances* of the two hypotheses are equal. In the second I have no reason to prefer one answer over the other as I have no reasons bearing on the matter at all. And of course there are various intermediate cases involving, say, partial knowledge of chances (perhaps I know only that between 40% and 60% of the marbles are black). One question I’ll be concerned with is what the epistemological significance of the difference between these cases is.

2. PRINCIPLE OF INDIFFERENCE

The Principle of Indifference links evidential symmetry and rational credence in a way that is blind to the difference between the cases above. Let $P(\cdot)$ be any *rational* subject’s credence function.

Principle of Indifference (POI): $p \approx q \rightarrow P(p) = P(q)$

An obvious corollary that is often called POI is:

POI:* If $\{p_1, p_2, \dots, p_n\}$ is a partition of your knowledge such that $p_1 \approx p_2 \approx \dots \approx p_n$, then for all i $P(p_i) = 1/n$.²

Let me set aside a common misunderstanding to begin with. One often hears: “You can’t get probabilities out of ignorance.”³ Let’s be clear that POI, as I am

² I’m taking for granted here that one’s credal state is best ideally represented by a single probability function mapping each proposition onto a real number. I’ll be questioning this assumption later.

³ As Michael Strevens (1998: 232) voices the concern: “It is surely the case that we can never get reliably from ignorance to truth... The fact that we do not know anything about A does constrain the way things are with A.”

understanding it, puts a *normative* constraint on what your *credence* may be. It entails that in a position of ignorance you are not rationally permitted to be more confident of one proposition than another. It is not to be confused with a principle for determining what the *objective probabilities* or *chances* are, where these are understood as supervening on the laws and physical properties of objects (e.g. that this coin has a 1/2 chance of landing heads on the next toss has to do with its shape, mass distribution, and manner of tossing; it is not a matter of my attitudes or evidence). Obviously ignorance is no basis for a *belief* concerning contingent physical conditions. But it is not at all out of the question that your ignorance puts constraints on what your degrees of confidence should be. I hope we agree that if I have no more reason to suppose that it will rain than that it won't then I should not be *absolutely certain* that it will rain, or even fairly certain. POI takes this idea further by insisting that if I am to be *any* more confident that it will rain than not, I had better have some *reason* for this difference of opinion.

Misunderstandings on this point may be due to people having meant different things by 'Principle of Indifference'. For instance, some defenders of POI (or at least of what they call 'POI') claim in its defense that the principle has often proved successful in predicting the behavior of stochastic processes.⁴ If we were in the business of determining objective chances then such success would indeed be relevant. That the frequency of heads in repeated coin flips matches the chance we assigned to it is evidence that our chance assignment was correct. But if our point is just to assign rational degrees of confidence often in complete ignorance of objective chances, then whatever match there may be between assigned probabilities and observed frequencies is of no relevance either way.

Distinguishing objective and subjective probabilities helps us clear up an objection to POI raised by Kyburg (1970) and others:

Suppose we know that this coin is biased, but have no idea, and no relevant evidence concerning which way it is biased. POI recommends assigning a probability of 1/2 to its landing *heads* on any toss, as we have no more or less reason to suppose that it will land *heads* than land *tails*. Now if the probability that it lands *heads* on each toss is 1/2, we should expect it to land *heads* about half the time in a long series of tosses. But of course we should expect no such thing. Given that the coin is biased we should predict that it lands predominantly on one side.

What has gone wrong here? The move from probabilities to frequencies here is best understood as an application of the Weak Law of Large Numbers. A simplified instance will do for our purposes:

Let $H = \{h_1, h_2, \dots\}$ be a set of propositions and $T_n = \#$ true members of $\{h_1, h_2, \dots, h_n\}$. If P is a standard probability function such that:

⁴ A classic example is Jaynes (1973).

1. *Identity*: For all i $P(h_i) = x$
2. *Independence*: $\{a_1, a_2, \dots, a_n\} \subseteq H \rightarrow P(\&i a_i) = \pi_i P(a_i)$

then for a sufficiently large n , $P(T_n|n \text{ approximately equals } x)$ is very high.⁵ Letting h_i – *the coin lands heads on the i th toss*, we get the familiar result that we should expect a coin to land *heads* with a frequency matching the probability of *heads* in the long run.

It is the Independence condition that is crucial here. It entails that the probability of *heads* remains the same conditional on any information regarding the outcome of other tosses. If our credence of x for heads on each toss is based on knowledge of the objective chance of *heads* then our credence assignment should meet the Independence condition. For our credence in h_i should match the known chance unless we have other evidence relevant to h_i , (information concerning the outcome of other tosses could only be relevant to h_i via relevance to the chance of h_i). This is why we should expect frequencies to approximately match known chances in the long run. But if our identically distributed credence over the h_i is not based on known chances then Independence need not (indeed should not) hold. For example, conditional on the coin's having landed *heads* repeatedly it is considerably more likely that it will land heads again, for this data should increase our estimate of the objective chance of *heads* on each toss. Without the Independence assumption the absurd conclusion that we should expect a biased coin to land *heads* half the time does not follow. We can perfectly coherently give credence 1/2 to the coin's landing heads on any toss, while giving very low credence to it's landing *heads* even close to half the time.

3. MULTIPLE PARTITIONS PROBLEM

The most famous objection to POI alleges that it leads to inconsistent conclusions.⁶ A probability space can be partitioned in different ways. If a proposition p is a member of two evidentially symmetric partitions S_1 and S_2 of different size, then POI gives inconsistent answers as to what your credence in p should be. The most compelling examples of the problem involve continuous parameters non-linearly related. Here is an example based on Bas van Fraassen's (1989) cube factory story.

Mystery Square. A mystery square is known only to be no more than *two* feet wide. Apart from this constraint, you have no relevant information concerning its dimensions. What is your credence that it is less than *one* foot wide?

It would appear that you have no more reason to suppose that the square is less than 1 foot wide than that it is more than 1 foot wide. But its *area* could be

⁵ More precisely: $P(T_n/n = x \pm \epsilon) \rightarrow 1$ as $n \rightarrow \infty$ for any $\epsilon > 0$.

⁶ See Gillies (2000), van Fraassen (1989).

anything from 0 to 4 square feet. Have we any more or less reason to suppose that it is less than 1 square foot, than that it is between 1 and 2 square feet, or between 2 and 3, or 3 and 4?

So we have two possible partitions:

$L_1: 0 \leq \text{length} < 1 \text{ ft.}$	$A_1: 0 \leq \text{area} < 1 \text{ sq. ft.}$
$L_2: 1 \leq \text{length} \leq 2 \text{ ft.}$	$A_2: 1 \leq \text{area} < 2 \text{ sq. ft.}$
	$A_3: 2 \leq \text{area} < 3 \text{ sq. ft.}$
	$A_4: 3 \leq \text{area} \leq 4 \text{ sq. ft.}$

There is no coherent way to assign probability evenly over both partitions. But POI seems to entail that we should. Setting the matter out so we can see what assumptions we are making:

(1) $L_1 \approx L_2$	Premise
(2) $A_1 \approx A_2 \approx A_3 \approx A_4$	Premise
(3) $P(L_1) = 1/2$	(POI)
(4) $P(A_1) = 1/4$	(POI)
(5) $P(L_1) = P(A_1)$	(L_1 and A_1 are equivalent)
<i>Contradiction.</i>	

We arrived at this contradiction by POI, so POI must be false.

That's not quite right. We also used two premises, (1) and (2). Why suppose these are true? Well, it's hard to see any reason to suppose that L_1 is true rather than L_2 , or to believe any A_i more than another. It appears that we have no reasons bearing on the matter at all. No information was given besides the maximum width.

Nevertheless, here is an argument to suggest that (1) and (2) are not both true. With the help of the following two principles that seem obviously correct, we can derive an absurd conclusion from (1) and (2) without any use of POI.

Transitivity: If $p \approx q$, and $q \approx r$, then $p \approx r$ ⁷

Equivalence: If p and q are known to be equivalent, then $p \approx q$.

(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) $A_2 \approx (A_2 \vee A_3 \vee A_4)$	(Transitivity)

⁷ In correspondence Elliot Sober and Branden Fitelson have independently objected to Transitivity by appeal to sorites cases. I'm not convinced, although the matter requires more thought and can't be addressed properly here. In any case, the question in this context is whether we really think that my minimal application here involves a violation of transitivity and that this accounts for the odd conclusion.

Consequence: We have *no less reason* to suppose that the area lies between 1 and 2 sq. ft. than to suppose that it lies between 1 and 4 sq. ft.

But this seems obviously wrong. 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 . If somehow we could rule out A_3 and A_4 , then ($A_2 \vee A_3 \vee A_4$) and A_2 might be epistemically on a par. But A_3 and A_4 could easily be true even if A_2 is not. So if I'm to believe only what is true, ($A_2 \vee A_3 \vee A_4$) is a safer bet. Since this odd conclusion follows from the premises without any use of POI, this casts doubt on the premises that were used to refute POI. (It is not hard to see how we might adapt this response to other versions of the Multiple Partitions objection.)

Actually we can derive an even worse conclusion from premises (1) and (2) with the help of the following principle.⁸

Symmetry Preservation (SP): If $p \approx q$ and r is known to be inconsistent with both p and with q , then $(p \vee r) \approx (q \vee r)$.

The principle seems obviously correct. It is hard to see how disjoining the very same proposition, which is inconsistent with each of two evidentially symmetric propositions, could break the symmetry. For example, suppose you have no more reason to suppose the marble in my hand is black than that it is white, or vice versa. Could you have any more reason to suppose that it is (white or red) than that it is (black or red)? Here is one way to think of it. Given the antecedent of SP, there are just two exclusive ways for $(p \vee r)$ to be true. Perhaps r is true (and p false), but if so then so is $(q \vee r)$. No asymmetry there. Alternatively, p might be true (and r false). But by assumption we have no more or less reason to suppose that q is true, which would make $(q \vee r)$ true. No asymmetry there either. (If r were consistent with either p or with q then there would be the further possibilities $(p \ \& \ r)$ and $(q \ \& \ r)$ and this may introduce an asymmetry, say if r supported p more than q . But SP does not apply in this case).

From symmetry premises (1) and (2) above we had derived the conclusion

$$(5) A_2 \approx (A_2 \vee A_3 \vee A_4)$$

Now with the help of SP we get

$$(6) (A_1 \vee A_2) \approx (A_1 \vee A_2 \vee A_3 \vee A_4)$$

We have no less reason to suppose that the square is no more than 2 square feet than that it is no more than 4! But we have *every* reason to suppose that it is no more than 4 square feet. We *know* that it is since it is no more than two feet wide. We have very little reason to suppose that it is no more than 2 square feet. It could easily be bigger. Again, the premises used to impugn POI lead independently to an absurd conclusion.

⁸ Thanks to Kenny Easwaran and Mike Titelbaum for pointing this out.

While I believe this knocks the wind out of the most popular refutation of POI, this rebuttal is unsatisfying in a certain way. One is apt to ask,

Alright, so what *should* my credence that the square is no more than a foot wide be, according to POI? If your argument is sound and so (1) and (2) are not both true, it follows that either we have more reason to believe one of $\{L_1, L_2\}$ over the other, or we have more reason to believe one of $\{A_1, A_2, A_3, A_4\}$ over others. So which is it, and what is this elusive reason?

Well, okay, so I don't really have an answer. Part of what is puzzling here stems from the temptation to think that my reasons or evidence must be transparent to me. The arguments above suggest that I have some reason to believe one proposition over another, yet it is hard to say what this reason might be, or even which proposition it supports. These reasons, if there are such, seem to be rather mysterious, accessible if at all only to enlightened souls. But it is tempting to think that reasons that are beyond my ken can't really function as reasons for *me*. These are murky waters that I can't wade through here. Suffice it to say that there are reasons to resist the temptation to think that your reasons or evidence must always be known to you (see Williamson, 2000, 2008).

That still leaves us with a different sort of complaint:

Perhaps your argument casts doubt on the claim that the Multiple Partitions Problem proves that POI is *inconsistent*. But it does at least appear to be *useless*. If we can't tell when propositions are evidentially symmetric because reasons to believe can be so elusive, then we can't apply POI to determine what our credence ought to be.

We might pause here to consider what we want out of a principle like POI. As stated, the principle gives a very general sufficient condition for being rationally required to have evenly distributed sharp credences. Even if the present complaint is warranted it does not immediately follow that the principle is incorrect or that we don't know that it is correct. If POI is correct then it is of no small significance to epistemology. One central concern of epistemology is to chart broadly the conditions of rational opinion. Such general principles might also tell us *why* certain attitudes are required of us.

But sometimes we hope for more. We want principles that might serve as *guides*. We would like rules that will advise us on which attitudes to take once we have ascertained that we meet certain conditions. When it comes to the mystery square and other such cases I suppose the POI is pretty useless if we are hoping to settle what our credence ought to be (even if its prescription is binding on us, whatever it is). But it by no means follows that it is always useless. There are Bayesian principles that I'm too dense to apply to some cases, either because they are too complex or too confusing. But that doesn't stop me usefully applying them in easy cases. For all that the Mystery Square and related cases suggest, there may be plenty of cases in which I can tell perfectly

well that various possibilities are evidentially symmetric. And the POI might usefully reveal both *that* and *why* my credence should be evenly divided.

We see here both the strength and the weakness of the POI as I have formulated it. I suspect that many who are hostile to POI view it as trying to do something clearly misguided: taking purely structural features of a space of possibilities as giving conditions of 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 empty. If all carvings are allowed we get inconsistency. If further criteria are imposed they often seem arbitrary and unmotivated (think of early Carnapian (1950) attempts to give syntactic criteria for probability assignments). The whole approach seems like a mug's game.

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. The difficulty that comes with this is that questions of just what my reasons are can in some cases be at least as difficult as the question of what my credence ought to be.

4. ARGUMENTS FOR POI

4.1. *Argument from Cases*

In many cases we clearly should assign credence $1/n$ to each of n alternatives. E.g. there are three doors. Behind one is a prize. What is your credence that the prize is behind door A (before any doors are opened)? In textbook cases like this we have no trouble answering $1/3$. Only POI can account for this.

Reply: While we might be required to distribute our credence evenly in the suggested cases, this is really because we have some implicit knowledge that is relevant to the probabilities, we have some kind of stochastic model that justifies this distribution. In such cases we should just apply the Principal Principle (Lewis, 1980) and set our credence to the known chances. There may be other cases in which we have no such knowledge of chances but are still tempted to think that an assignment of equal credence is called for. We should resist this temptation, which might be based on a confusion with cases of the legitimate application of known chances.

As you can imagine, tossing back and forth cases like this quickly leads to a stalemate. But here is a case to think about that is based on a largely uncontroversial part of the Sleeping Beauty problem (Elga, 2000).

Sleeping Beauty Case. You know that you will be awakened on Monday and again on Tuesday, but your memory of the first awakening will be erased

before the second. Now you find yourself awake and unable to tell whether it is Monday or Tuesday. How do you divide your credence?

Although just about everything that could be said has been said about the original Sleeping Beauty problem and the broader issue of *de se* credence and updating, as far as I know everyone thinks that in the case above your credence should be divided evenly. (We can make the case more extreme: you are to be awakened every day for a year with no memory of previous awakenings. Surely you should put *low* credence in today being November 24.) But in this case it is hard to make sense of some kind of stochastic model at work. There is no process whose outcome determines whether *today* is Monday or Tuesday. There is no chancy event like a coin flip such that if it turns out one way it will now be Monday and if it turns out the other it will now be Tuesday. If we can speak tenselessly, you *are* awake on Monday *and* awake on Tuesday (and you know it). Your uncertainty is just about which day it is *now*. So it is hard to see how we can be appealing to any kind of implicit stochastic information in giving 1/2 credence to it's being Monday. Rather it seems just to be matter of our ignorance concerning which day it is. You have no reason to suppose that it is one day rather than the other. It is hard to see how to defend the 1/2 answer without appeal to POI.

4.2. *Argument from Statistical Inference*

Many who are hostile to POI see at as a spooky a priori thing where you mysteriously conjure probabilities out of pure ignorance. Many people who feel this way will contrast it with the case of basing one's credence on *known statistics*, good solid data concerning the frequency of certain types of event or attributes in a population. For example, you somehow learn that 37% of formal epistemologists are left-handed. With *only* this to go on, what should your credence be that Branden Fitelson is left-handed? The very natural answer is 37%. Of course the matter gets tricky if I also know that, say, 17% of University of Wisconsin Madison graduates are left-handed, but have no direct information about the frequency of left-handedness among formal epistemologists from Madison. Or perhaps I saw him catch a ball twice with his right hand and twice with his left. But provided we have none of this kind of messy conflicting data, and nothing more specific to go on, setting one's credence to known frequencies seems clearly correct. The general principle at work here, which sometimes goes by the names Statistical Syllogism or Direct Inference, is something like the following.

Frequency-Credence (FC). If (i) I know that a is an F , (ii) I know that $\text{freq}(G | F) = x$ (the proportion of F s that are G), and (iii) I have no further evidence bearing on whether a is a G , then $P(a \text{ is a } G) = x$.

Perhaps there is a more general principle covering cases where I know that a is also an H and that $\text{freq}(G | H) = y \neq x$. It is notoriously difficult to come up with

a satisfactory one.⁹ Still, something like the more restricted FC seems straightforward. Quite a number of philosophers who express doubts or even hostility to POI seem to endorse something like FC.¹⁰ The curious thing is that FC entails POI*.

Proof: Let $F = \{p_1, p_2, \dots, p_n\}$ be any set of disjoint and exhaustive possibilities such that $p_1 \approx p_2 \approx \dots \approx p_n$. Let G be the set of *true* propositions. For any p_i : (i) I know that p_i is an F (i.e. that $p_i \in \{p_1, p_2, \dots, p_n\}$); (ii) I know that $\text{freq}(G | F) = 1/n$ (exactly one member of the partition $\{p_1, p_2, \dots, p_n\}$ is true); and (iii) I have no further evidence bearing on whether p_i is G (I am ignorant concerning the p_i , with no more reason to suppose that one is true rather than another). Hence by FC, $P(p_i \text{ is a } G) = 1/n$, i.e. $P(p_i \text{ is true}) = 1/n$, so $P(p_i) = 1/n$.

It's an illusion that POI involves magically getting knowledge out of ignorance while FC solidly grounds probability judgments in data. Of course this might cast more doubt on FC than support POI. But since we obviously should use frequency data to guide our subjective probabilities in *something* like the FC way, the opponent of POI faces the challenge of identifying the correct principle in this vicinity.

It's worth noting here the parallel between the Multiple Partitions problem and a version of the Reference Class problem. Just as a proposition may be known to be a member of different partitions of a space of possibilities, an individual like Branden can be known to be a member of different classes. Naive applications of either POI or FC can lead to inconsistent results. And we lack an adequate story concerning how to extend either principle to the tricky cases. But now while this is a difficult problem, it is seldom taken to show that frequency data are *never* a legitimate guide to credences. Should the Multiple Partitions problem be taken to show that POI has *no* legitimate application?

Nevertheless, quite independently of worries about POI, some philosophers such as Isaac Levi (1977) have objected to FC, saying that we are required to match credence to frequency only in cases where we have *randomly selected* the item in question from the population. It is not just that we must not have reason to suppose that the way in which we came to consider Branden as our example was biased in favor of (or against) left-handers. We must know that the process by which he was selected was an objectively random one, i.e. each individual in the reference class had equal objective chance of being selected. According to Levi's account, learning that Marble A is among a hundred marbles in a certain urn, 37 of which are black, is of little relevance in assigning credence to A's being black. But suppose we shake the urn and select a marble in a suitably random manner. Then prior to seeing its color our credence that it is black should be 37%.

⁹ A valiant attempt is made by Kyburg (1974).

¹⁰ For example, Hacking (1965), Hacking (2000), Hajek (2007), Kyburg (1974), and McGrew (2003).

I find it hard to reconcile the common lines of statistical reasoning we all engage in with Levi's position. It is a little unclear how the appropriate kind of 'random selection' is meant to work in less contrived cases. I ride a motorcycle and know something of the accident statistics for riders in different classes. Surely this sort of information should inform my credence in the possibility of a crash. (The insurance companies are certainly using it!) What would it mean for me to be randomly selected in the relevant way? It is not as if I picked someone at random out of the directory of motorcyclists and it happened to be me. I started with myself and went looking for statistical data that might apply to me. So I have trouble seeing how this random selection condition on FC is supposed to be applied if it is to do justice to our actual inferential practices.¹¹

But in any event, if we really need to do this random sampling then there is nothing to stop us doing it. Let's take all the formal epistemologists and toss them into a big urn. We will shake it vigorously and pick someone out. It happens to be Branden. Are we now supposed to have 37% credence that he is left-handed? If so then even if we don't get Branden the first time, we could just keep randomly selecting until we do. While this might be fun, it doesn't seem to be necessary.¹²

4.3. Evidentialist Argument

One's confidence should adequately reflect one's evidence (or lack of it). You need a good reason to give more credence to p than to q . Hence if one's evidence is symmetrical so should be one's degrees of confidence. This is the fundamental thought behind POI and it can't be easily dismissed. What other option is there? Could there be some other rule determining what your credence should be in cases of evidential symmetry, particularly in those

¹¹ The motorcycle case might not be so straightforward as the frequency information might be relevant to estimates of *chances* of accidents in various situations, and these chances are then a guide to my credence. Perhaps most ordinary cases have an element of this. But I don't think that's essential. Suppose an oracle who is known to see the future reveals that 99% of motorcyclists will have an accident this summer. She assures me that this is *not* because the conditions will be any more dangerous than usual. Each rider will have about the same objective chance of an accident as is usually the case. There will just be a fluky series of unlikely accidents (much as a perfectly fair coin will occasionally have a long run of heads). Shouldn't I have a high expectation of an accident this summer (given, of course, that this won't stop me riding)?

¹² One of Levi's complaints about FC is that (Kyburg's version of it) conflicts with the rule of conditionalization. I don't think my statement of FC faces this problem. I can't get into the details of Levi's argument here, but there is an interesting parallel with what I earlier mentioned as the strength and weakness of my version of POI. Kyburg attempts to give criteria for probability assignments purely from frequency information. The appeal of this is that it is straightforward enough to collect such information and then we can just read off what the correct credence is that an individual has a certain attribute. My statement of FC contains the crucial *epistemic* clause: (iii) I have no further evidence bearing on whether a is a G . While I believe this avoids Levi's objection, it makes the principle less straightforwardly applicable as it may often be a murky matter whether condition (iii) is met.

cases where the symmetry is due to ignorance? Perhaps the rule is that one's credence must be divided in the ratio .327 : .673 among two possibilities. But this is silly. To even be applicable we must specify some difference between the two propositions to determine which one gets to enjoy the .673 and which the .327. And the only factor that could sensibly play such a role in determining these credence assignments would be some kind of *evidence* or *reasons* in support of one proposition (in which case we do not have evidential symmetry). Alternatively we might hold a *permissive* view, according to which no particular credence distribution is rationally required; any of some range is a rational option.¹³ But how wide is this range? How about 1:0? No, it would be nutty to be *certain* that *p* rather than *q* on the basis of no relevant evidence. (Talk about getting knowledge magically out of ignorance!) Would .9 : .1 be okay? That's just about as bad for the same reason. Smaller divergences from a nice .5 : .5 might not be *as* crazy as being highly confident that *p* rather than *q* for no reason. But the same scruples that prevent us from the more drastic imbalances of opinion reveal that ideally we should just split our credence evenly unless there is a reason to do otherwise.

Perhaps examining a case of practical importance will focus our judgment. A homicide detective has found only two suspects, Alice and Bob, who could have committed the murder. He reveals that at this stage he is at least somewhat more inclined to think that Alice is the guilty one. "Why?" We ask. "No reason," he replies. "Unfortunately at this stage I have no evidence to go on." "But you have some kind of reasonable hunch, right? You must take yourself to have some kind of rational insight that points toward Alice, even if you can't articulate it as an item of evidence." "Oh no, nothing of the sort. I have nothing at all that could count as any kind of reason to suppose that it was Alice rather than Bob. I just happen to be more confident that it was Alice." This can seem downright perverse. He is trying to get at the truth here and he just biases the investigation one way for no reason! If he updates his confidence in a reasonable way as he gathers more evidence he will reach the level of confidence required to convict Alice on the basis of less evidence than if he had begun his investigation being neutral with regard to the suspects. In court he will have to admit that part of the basis for his confidence that Alice is guilty is just an arbitrary bias he holds for no reason at all. Or perhaps the only evidence he finds moderately supports the theory that Bob is the murderer. The evidence is just strong enough to even up the detective's credence in his two hypotheses. Although the only evidence he has suggests that *Bob* did it, he remains blithely neutral on the matter!

Now I take it that the fact that justice is at stake here is of no *epistemological* relevance. It just helps to sharpen our focus on the issue. A real detective who proceeds in this manner should be taken off the case. We can't trust him to reason in a manner that will most likely lead to the right conclusion. We are

¹³ North (MS) suggests this. I discuss permissive epistemologies more generally in my 2005 paper.

less concerned with the 'armchair detective' who deliberates in a similar way while watching the case on TV. But it is hard to see how he can be any less rational given that he is in the same evidential situation. It is just that it is more important that a real detective be epistemically rational. And in any such case it appears that the only appropriate criticism is that the subject is assigning disproportionate credence without any difference in evidence. That is, he is violating POI.

Reply: There is something right and something wrong about this line of thought. Yes, evidential symmetry demands symmetry of opinion. But by failing to distinguish cases of *known chances* and cases of *ignorance*, the follower of POI actually *fails* to adequately represent his evidential state. The mistake is to suppose that one's opinion must always be represented by the simple model of a single standard probability function. Only in cases of known equal chances should one's credence be divided sharply in this manner. Ignorance calls for a different kind of state in which one's credence is spread, as it were, over a range of values.¹⁴

I find this the most compelling alternative to the standard POI. The idea that one's credence should in some sense cover a range of values seems to be orthodoxy these days. The phenomenon is given a number of different names, such as 'indefinite' credence (Joyce, 2005), 'vague probability' (van Fraassen, 1990), 'indeterminate probability' (Levi, 1974) 'imprecise probability' (Walley, 1991), 'thick confidence' (Sturgeon, 2008), and others. I've been following Elga (MS) in calling it 'mushy' credence (although the novelty of this terminology is wearing off). It is perhaps not entirely clear that everyone has the same idea in mind. But there is at least a common kind of formal model whose main elements are these:

Representors. A rational subject's state of opinion is best represented by a *set* of probability functions called his *representor R*.

Updating. Each of the functions in the representor is to be updated by conditionalization on new evidence.

Often we are just interested in the spread of values in our representor for a particular proposition, i.e. the range of values for which there is some function in one's representor assigning that probability to the proposition. I will speak of one's *credence* in a proposition as being possibly a set of values. To avoid confusion I'll use '*P*' for a standard probability function mapping propositions onto single real numbers, and '*C*' for one's credence function as follows: $C(p) = \{x : \exists P \in R, P(p) = x\}$.

The key idea behind this response to the evidentialist argument is the following:

¹⁴ Joyce (2005) best articulates this reply.

Chance Grounding Thesis. Only on the basis of known chances can one legitimately have sharp credences. Otherwise one's spread of credence should cover the range of possible chance hypotheses left open by your evidence.

We can illustrate the idea with the following case from Joyce (2005). I have before me three urns that you know to be taken from a collection $\{urn_0, urn_1, \dots, urn_{10}\}$ such that urn_i contains i black balls and $10 - i$ white balls. Here is what you know about each urn:

U_1 : You know only that this is urn_5 .

U_2 : You know only that this urn was chosen at random from the eleven urns.

U_3 : You know nothing about which urn this one is, or how it was chosen.

In each case, what credence should you have that a random selection from the urn yields a *black* ball? For U_1 the answer is clearly $1/2$. According to Joyce, since you have no relevant chance information concerning U_3 , your credence that it will yield a black ball should be $[0,1]$. As for U_2 , you don't know how many black and white balls it contains, so you don't know this urn's chance of yielding a black ball. But you have some higher-order chance information. You know that U_2 was selected from the urns such that each urn had the same chance of being selected. So you should have sharp credence $C(U_2 = urn_i) = 1/11$ for all i . And sharp conditional credence $C(U_2 \text{ yields a black ball} \mid U_2 = urn_i) = i$. So your credence that U_2 will yield a black ball should be $C(U_2 \text{ yields a black ball}) = \sum_i C(U_2 \text{ yields a black ball} \mid U_2 = urn_i) C(U_2 = urn_i) = 1/2$.

Joyce's example illustrates the Chance Grounding Thesis. But it also raises a puzzle related to our earlier discussion of random selection. Suppose all eleven urns are lined up in a row in an unknown order. You have no idea by what method they were arranged. With respect to each of these urns you would appear to be in the same situation as with U_3 above. You have no relevant information concerning chances. So on Joyce's account it seems that your credence for each urn that it will yield a black ball should be mushy over the range $[0,1]$. I might point to each urn in order along the row asking "What's your credence that this one will yield a black ball?" And in each case you shrug your shoulders, unable to pin it down to any range narrower than $[0,1]$.

But now I start jumping about and flinging my arms around wildly over the urns. You happen to know that my arms are tossing about in an objectively random manner, so that there is an equal chance of my hand landing on any particular urn at a time. Each time I touch an urn I ask once again, "What's your credence that this one will yield a black ball?" Now your situation seems similar to U_2 above. You know that this urn to which I direct your attention has been selected at random from the eleven. So on Joyce's account your credence that it will yield a black ball should apparently be sharply $1/2$. Somehow by magically waving my hands over the urns I have sharpened up your credence dramatically. This can't be right.

5. THE COIN PUZZLE

The following case illustrates something extremely puzzling about the mushy credence picture as a response to evidential symmetry.

Coin game. You haven't a clue as to whether p . But you know that I know whether p . I agree to write ' p ' on one side of a fair coin, and ' p' ' on the other, with whichever one is true going on the heads side (I paint over the coin so that you can't see which sides are *heads* and *tails*). We toss the coin and observe that it happens to land on ' p' '.

Let C and C_+ be your rational credence functions before and after you see the coin land, respectively. The following five propositions are jointly inconsistent:

- (1) $C(p) = [0, 1]$
- (2) $C(\text{heads}) = \{1/2\}$
- (3) $C_+(p) = C_+(\text{heads})$
- (4) $C_+(p) = C(p)$
- (5) $C_+(\text{heads}) = C(\text{heads})$

I'll consider each in turn.

- (1) $C(p) = [0,1]$ Prior to seeing the coin land your confidence in p should be spread over the entire $[0,1]$ interval in a maximally agnostic state.

According to the mushy credence response to POI, this is the attitude one should take toward a proposition in a position of complete ignorance. It is the most noncommittal attitude one can take (or what I've heard described as 'the ultimate shrug'). Let's suppose this is so for *reductio*. I think the problems raised here arise for less extreme mushiness also. But for simplicity we can focus on the extreme case.

- (2) $C(\text{heads}) = \{1/2\}$ Prior to seeing the coin land, your confidence that it will land *heads* should be sharply $1/2$. In other words, *every* function in your representor should assign $1/2$ to *heads*.

Your credence should be set to the known objective chance of *heads*. This is an uncontroversial application of what Lewis (1980) called the Principal Principle. And at any rate this principle is just intended to accommodate obvious facts like (2).

- (3) $C_+(p) = C_+(\text{heads})$ Upon seeing that the coin lands ' p' ', your confidence in the propositions p and *heads* should be the same.

You know that ' p' ' is on the *heads* side iff p is true (I wrote the true proposition on the *heads* side), and that it landed on ' p' '. So you know that it landed *heads* iff p . Hence your attitude to these propositions should be the same.

If we accept the assumptions thus far, some change in my attitude to either p or to *heads* is required upon seeing the coin land on ' p' '. Either my credence in

p should *sharpen* to match *heads* at $1/2$, or my credence in *heads* should *dilate* to match p at $[0,1]$.¹⁵ The last two premises say that no such change is appropriate.

(4) $C_{-}(p) = C(p)$ Seeing the coin land ' p ' should have no effect on your credence in p .

We might deny this and say that upon seeing the coin land ' p ' you should sharpen your credence in p to $1/2$. But note that if we say this then for symmetrical reasons we will have to say the same if the coin lands ' $\neg p$ ' instead. For if the coin lands ' $\neg p$ ' you will then know that p is true iff the coin lands *tails*, and your credence in *tails* will be $1/2$. On this view you will know in advance that your credence in p will be $1/2$ *no matter how the coin lands*. But this can't be right. If you really know this in advance of the toss, why should you wait for the toss in order to set your credence in p to $1/2$?

I won't spend so much time addressing the possible denial of (4) as many of the arguments I make in the next section for (5) can be modified to apply to (4). But more importantly, the suggestion that one should sharpen one's credence to $1/2$ in this case is incompatible with the standard rule of updating each function in one's representor by conditionalization. The following is a theorem of probability.

Irrelevance. For any probability function P , $P(p | e) - P(p | \neg e) \rightarrow P(p | e) = P(p)$

If I am to update by conditionalization such that all of the functions in my representor converge to $1/2$ when I see the coin land (whether it lands ' p ' or ' $\neg p$ '), then prior to the toss each function must be such that

(*) $P(p | \text{coin lands } 'p') = P(p | \text{-coin lands } 'p') = 1/2$

But by Irrelevance it follows that for each P in my representor, $P(p) = 1/2$, which contradicts our assumption (1) $C(p) = [0,1]$. (Indeed this is just the conclusion recommended by POI.)

(5) $C_{+}(\text{heads}) = C(\text{heads})$ Seeing the coin land ' p ' should have no effect on your credence in *heads*.

Our last option to save (1) is to deny (5) and say that your credence in *heads* should *dilate* to $[0,1]$. As before, if we take this option, then for symmetrical reasons we must say the same if the coin lands ' $\neg p$ '. This option is actually mandated by the standard account of updating. For when we learn that the coin landed ' p ' we learn that $(p \leftrightarrow \text{heads})$; if the coin lands ' $\neg p$ ' then we learn $(\neg p \leftrightarrow \text{heads})$. Now, each function in my representor must be such that

(**) $P(\text{heads} | p \leftrightarrow \text{heads}) = P(p)$, and

(***) $P(\text{heads} | \neg p \leftrightarrow \text{heads}) = 1 - P(p)$

¹⁵ I suppose someone could say that my credence in both p and *heads* should change and meet in the middle at some narrower range of values. I don't think this suggestion requires separate consideration.

Proof of ():**

$$\begin{aligned}
P(\text{heads} | p \leftrightarrow \text{heads}) &= P[\text{heads} \ \& \ (p \leftrightarrow \text{heads})] / P(p \leftrightarrow \text{heads}) \\
&= P(p \ \& \ \text{heads}) / P[(p \ \& \ \text{heads}) \vee (\neg p \ \& \ \neg \text{heads})] \\
&= P(p \ \& \ \text{heads}) / [P(p \ \& \ \text{heads}) + P(\neg p \ \& \ \neg \text{heads})] \\
&= P(p)P(\text{heads}) / [P(p)P(\text{heads}) + P(\neg p)P(\neg \text{heads})] \\
&= P(p) / [P(p) + P(\neg p)] \\
&= P(p)
\end{aligned}$$

(Similar proof for (***))

According to (1), for all $x \in [0,1]$, I have a function P in my representor such that $P(p) = x$. It follows now from (**) and (***) that when I update on either $(p \leftrightarrow \text{heads})$ or $(\neg p \leftrightarrow \text{heads})$ (i.e. on the coin's landing ' p ' or its landing ' $\neg p$ ') I will have a P in my representor such that $P(p) = x$, for all $x \in [0,1]$. That is, my credence in *heads* dilates from $C(\text{heads}) = \{1/2\}$ to $C_+(\text{heads}) = [0,1]$.

Having noodled about this puzzle on and off for some time, I discovered that the general phenomenon of dilation is old news.¹⁶ Some statisticians and philosophers have studied how the phenomenon arises in other cases and appear to have taken it in their stride. This is not a *reductio* but a *result*, they might say.¹⁷ I want to suggest that the present case brings out particularly forcefully how bizarre this phenomenon is, at least in the present case where we are assuming evidential symmetry between p and $\neg p$. Here is a series of objections to the view that one should dilate one's credence in *heads* to $[0,1]$ upon seeing the coin land one way or another.

Objection 1. Known Chance: You still know that the coin is fair. You have lost no information. How could the information you've gained have any relevance to your attitude to *heads*? Lewis (1980) was careful to point out that once you gain some 'inadmissible' evidence concerning the outcome of a fair coin toss—e.g. by seeing it land or receiving a prophecy about the outcome before it is tossed—you might legitimately have credence other than $1/2$ in the coin's landing *heads*. Suppose that in the present case instead of being clueless as to whether p you had some reason to suppose that p . In this case upon seeing the coin land ' p ' you would have gained some evidence that the coin landed *heads*. For now your evidence suggests that the coin landed on the side with the *true* proposition, and you know that the true proposition is on the *heads* side. (Note that in this case we don't get the same result if the coin lands ' $\neg p$ '; if the coin lands ' $\neg p$ ' you've gained some evidence that the coin did *not* land *heads*). So it is not out of the question that seeing the coin land ' p ' could be relevant to your credence in *heads*. But in the original case *you haven't a clue as to whether p*. You have *nothing* to suggest that the coin landed *heads* or that it landed *tails*.

¹⁶ See e.g. Seidenfeld and Wasserman (1993) and Walley (1991).

¹⁷ Others, like van Fraassen (2005, 2006), find the phenomenon more disturbing.

Shouldn't you just ignore this useless bit of information and keep your credence in *heads* at $1/2$?

Objection 2. Reflection: It is natural to suppose that if you know that you will soon take doxastic attitude *A* to *heads* as a result of *rationally* responding to new information *without loss of information*, then you should *now* take attitude *A* to *heads*. (This is a generalization of Bas van Fraassen's (1984) Reflection principle.) Why would you need to wait until you actually gained this information to adopt this attitude? It would be different if you expected to lose your marbles tomorrow and have poor judgment, or perhaps forget some important fact that would make a difference to your conclusion (Talbot, 1991). But as long you know that you will be epistemically fine then you should trust your future judgment and match your current attitude to it. But now according to the current proposal you know in advance that once you see the coin land you will rationally dilate your credence in *heads*. So you ought to just dilate it now before the coin even lands. But this is to deny (2), which is absurd.

We can make a similar point considering the opinions of another. Suppose you know that Scott takes attitude *A* to *heads*. You also take him to (a) have exactly the same 'priors' as you, (b) be impeccably rational, and (c) to possess *all of your knowledge plus more*. Surely you should trust his judgment and adopt attitude *A* yourself. That is after all the attitude that *you* would take if you were in his shoes, and he has the epistemic advantage of having more information to go on. This is just an easy case of how we trust experts more generally. Given (a)—(c), that Scott takes attitude *A* entails that he has information that warrants attitude *A* by *your* lights, so you should follow him. But now suppose the coin has been tossed and you know that Scott has seen it land while you have not. According to the current proposal, without even having to ask him you know that Scott's credence in *heads* is now mushy (for he is as ignorant as to whether *p* as you). So now yours should be mushy too. We have reached the absurd conclusion that your credence that this fair coin lands *heads* changes from $1/2$, not by receiving any information about how the coin landed (not even whether it landed '*p*') but just by learning that a rational onlooker has learned whether it landed '*p*'.

Objection 3. Mushy Betting: The state of having credence $C(p) = [0,1]$ can seem rather similar to $C(p) = \{1/2\}$. In each case you are no more inclined to suppose that *p* than to suppose that $\neg p$. If there is an important difference between the two states of opinion then presumably it will be manifest in the behavior of a rational agent. If your credence in *p* is $[x, y]$, at what odds should you bet? There are many possible answers, but here are the two most common answers I've heard:

Liberal. Take any set of bets that maximizes expected utility according to *some* credence function in your representor.

Conservative. Only bets up to $x : 1-x$ on p , or up to $1-y : y$ on $\neg p$ are permissible. Decline a bet offered in between if you can.

These prescriptions are different enough that I suspect they reveal rather different understandings of what it is to be in a mushy credal state. But let's consider them in turn.

Objection 3.1. Liberal betting: We are supposing that $C_+(heads) = [0,1]$. So upon seeing the coin land one way or another, among other things you have a function P in your representor such that $P(heads) = 3/4$. At this point someone offers you a bet at 2:1 on *heads*, (i.e. if it landed *heads* you win \$1; if not, you lose \$2). According to the Liberal account you can, as it were, plump for $P(heads) = 3/4$ as your credence for betting purposes. And you maximize utility according to this function by taking the 2:1 bet. So when you see the coin land you can bet at odds 2:1 on *heads*.

But that seems mad. Let's repeat the whole game over and over using a series of propositions p_1, p_2, \dots , such that $C(p_i) = [0,1]$. On each toss when you see the coin land either ' p_i ' or ' $\neg p_i$ ' you will dilate your credence that the coin landed *heads* on that toss to $C_+(heads_i) = [0,1]$. And by the Liberal account you will be permitted to bet at 2:1 on *heads_i*. I think we know what will (almost certainly) happen when you do this. This *fair* coin we are using will land *heads* about half the time and sooner or later you will go broke.

Here is another Liberal betting worry that is related to Reflection. Suppose you are offered a 2:1 bet on *heads* prior to the coin toss. According to the dilation story, prior to seeing the coin land you know that you will soon have credence $C_-(heads) = [0,1]$ when you see the coin land. And if you accept Liberal betting, then you believe that it will soon be rational for you to take 2:1 bets on *heads* (if offered only that bet or nothing). But in that case prior to the coin toss you can rationally adopt a policy to accept such a bet once you have seen the coin land. What could be wrong with a policy to accept a bet when it is rational to do so? But adopting such a policy is the same as just betting on *heads* at 2:1 now, *prior to the toss*. But that's mad. Your credence in *heads* prior to the toss is sharply 1/2. You can't sensibly take a bet at 2:1 in that case.

Objection 3.2. Conservative betting: The conservative betting policy will not allow foolish series of bets like those above. But it rules them out at the cost of disallowing bets that are obviously wise. Just switch the case around. On each toss you are offered a bet at 1:2 on *heads_i*; once you see the coin land ' p_i ' or ' $\neg p_i$ '. Since your credence in *heads_i* is mushy at this point you turn down all such bets. Meanwhile Sarah is looking on but makes a point of covering her eyes when the coin is tossed. Since she doesn't learn whether the coin landed ' p_i ' her credence in *heads_i* remains sharply 1/2 and so takes every bet. (Of course she does. At 1:2 the odds are strongly in her favor.) Sure enough, she makes a killing. "Don't you want to get in on this?" she asks. "I can't," you reply. "I keep seeing how the coin lands, so none of these bets are rational for me." Eventually you cave and join Sarah in making money hand over fist. But of

course you don't bother to close your eyes. It is not as though Sarah's lack of knowledge concerning the ' p_i ' and ' $\neg p_i$ ' outcomes has anything to do with her success. The coin is going to land *heads* about half the time regardless of whether anyone is watching.¹⁸

As before, there is a further problem for Conservative betting related Reflection. Prior to seeing the coin land you are offered a bet at 1:2 on *heads*. You are also offered the option of calling off the bet once you see it land ' p ' or ' $\neg p$ '. According to the dilation account and Conservative betting strategy, you know that once you see the coin land you will want to renege on the bet. For you know that your credence in *heads* will dilate and the Conservative betting cautions against bets within your mushy range of credence. Indeed if you accept these you should have a policy of canceling the bet once you see the coin land. But to have such a policy and know that you will cancel the bet really amounts to rejecting the bet even before you see the coin land. But you don't want to do that! If your credence in *heads* prior to the toss is sharply $1/2$ —as of course it should be—then 1:2 odds are strongly in your favor. Taking such a bet ought to be a no-brainer.

Objection 4. Many Coins again: The objections involving repeated tossing from the previous section depend on particular decision theories for mushy credence. Perhaps there are other options for betting that avoid these worries.¹⁹ But I think the basic worries can be brought out independently of decision theory. Suppose we repeat the experiment as above using a series of propositions p_1, p_2, \dots , such that $C(p_i) = [0,1]$. Upon seeing the coin land ' p_i ' or ' $\neg p_i$ ' for each toss we are supposed to have mushy credence in *heads_i* for that toss. Suppose that on each toss we also get to remove the labels and see if the coin did land *heads*. It will be hard not to notice after a while that *about half the time* when you are in the mushy state concerning *heads_i* the coin does land *heads*. And it will be hard not to infer *inductively* that it will continue to be the case that about half of the occasions in this scenario in which you've seen the coin land and your credence is $C(\text{heads}_i) = [0,1]$, *heads_i* is true. Now if you know that about half of the time when you are in a certain evidential situation it turns out to be raining, then surely when you next find yourself in such an evidential situation your credence that it is raining should be very close to $1/2$. So similarly in the case as described your credence in *heads_i* should be $1/2$ when you see the coin land ' p_i ' or ' $\neg p_i$ '. But you didn't really need to learn this inductively, did you? Wasn't it obvious from the beginning that this fair coin would land *heads* about half the time regardless of whether you have seen it land?

¹⁸ On a variant of conservative betting you are rationally permitted to either take or leave bets corresponding to values in the mushy region of your credence (provided the combination of your bets is not a Dutch book). This faces the same problem as it would be foolish *not* to bet in the present case.

¹⁹ If the answer is that one should bet *as if* one has sharp credence we might wonder once again what the difference between sharp and mushy credence really consists in.

Here is a similar puzzle. I label and toss one million fair coins with propositions as before. This time suppose that your credence in the proposition is only *slightly* mushy, say $C(p_i) = [0.4, 0.6]$. And suppose that the p_i are *independent* across all the functions in my representor. (Perhaps $p_i =$ 'a black ball will be selected from Urn $_i$ ' where each urn contains ten marbles, between four to six of which are black, and they have each been filled in a manner causally independent of the others).

Prior to reading how any of the coins landed, what is your credence in the following proposition?

half-heads: About half the coins—say 45 to 55% of them—landed *heads*.

You know it is overwhelmingly likely that about half of a group of a million fair coins will land *heads*, so your credence in *half-heads* is very high. What about your credence in the following?

half-true: About half of the propositions $p_1, p_2, \dots, p_{1000000}$ —say 45 to 55% of them—are true.

You have a function P' in your representor such that for each i $P'(p_i) = 0.6$. And another $P''(p_i)$ such that for each i $P''(p_i) = 0.5$. (And there will be plenty of functions in between, assigning different values to different p_i). Now,

P' (*half-true*) must be very low.

P'' (*half-true*) must be very high.

So $C(\text{half-true}) = [x, y]$, where $x \approx 0$ and $y \approx 1$.

Now according to the dilation account, when you read how each coin landed, for all i , $C_+(heads_i) = C(p_i)$. And so $C_+(\text{half-heads}) - C(\text{half-true}) = [x, y]$, where $x \approx 0$ and $y \approx 1$. That is, your credence that half the coins landed *heads* should now dilate to a highly agnostic state covering almost the whole range from 0 to 1. But now ask yourself seriously if this is what you really think. You are looking at pile of a million *evenly weighted coins, each of which has a track record of landing heads about as often as tails*, labeled with a bunch of propositions ($p_1, p_2, p_3, p_4, \dots$) about which you have no more evidence one way or another. Do you really doubt that about half of them landed *heads*?

We can drive this worry further by supposing instead that for all i , $C(p_i) = [0, 1]$. Now by reasoning along the lines above, the dilation account entails that $C_+(\text{all the coins landed heads}) = [0, 1]$. That is, upon reading for each coin whether it landed ' p_i ' or ' $\neg p_i$ ' you will become maximally agnostic concerning whether *every one of a million fair coins landed heads*. I doubt that you are. I'll bet you're about as confident as you are about anything that they *didn't* all land *heads*.

6. A FEW MORE PUZZLES

Here are some more problems for the mushy credence approach to representing symmetrical ignorance that also serve as arguments for POI. There is a

common theme in each, but I would like to drive the problem home by getting at it in different ways.

6.1. Coin Puzzle*

When I first started shopping the Coin Puzzle around I had a different version:

*Coin Game**. None of us have any clue whether p (we have no more reason to suppose that p than to suppose that $\neg p$ or vice versa). We stick ' p ' and ' $\neg p$ ' labels on either side of a fair coin (which sides they go on is irrelevant). We now toss the coin and observe whether it lands ' p ' or ' $\neg p$ '.²⁰

Prior to tossing the coin, our credence in the proposition

lands true: the coin lands on the side with the true proposition

should surely be $1/2$. Only one side has a true proposition and whichever side it is the coin has a $1/2$ chance of landing on it. How about when the coin lands and we read it? That would appear to have no relevance to whether the coin landed on the true side (it would if we had any evidence for or against p). So we have the following:

$$C(\textit{lands true}) = C_+(\textit{lands true}) = 1/2$$

But now upon seeing the coin land ' p ', every function P_{\perp} in my representor is such that

$$\begin{aligned} P_{\perp}(\textit{lands true}) &= P(\textit{lands true} \mid \textit{lands 'p'}) \\ &= P(p) \end{aligned}$$

Hence $C(p) = C_+(\textit{lands true}) = 1/2$.

Here we have neat little proof of POI with respect to any p and $\neg p$! The mushy credence opponent of POI will want to get off the boat before it leaves and deny that $C_{\perp}(\textit{lands true}) = 1/2$. As with *heads* in the previous puzzle, he will have to say that my credence in *lands true* should dilate to match p regardless of which way the coin lands.

This version of the puzzle has both dialectical strengths and weaknesses over the previous one. On the one hand it is somewhat simpler and I have found that some audiences have been suspicious about the way the agent who labels the coins in the previous puzzle knows whether p and uses this knowledge to the true proposition on the *heads* side. (I haven't heard a clear argument as to why this should undermine the point of the puzzle.) On the other hand, I found that when I presented the current version opponents of

²⁰ David Christensen and John Burgess also independently suggested Coin Puzzle* as a variation on Coin Puzzle. Scott Sturgeon considers yet another variant in his contribution to this volume (Ch. 5).

POI were more willing to say that dilation of credence in *lands true* was called for upon seeing the coin land. (It takes more guts to say that your credence that *a fair coin landed heads* should be anything but 1/2). Nevertheless, I think that most of the same worries can be raised for the dilation of *lands heads*. For instance, if we repeat the experiment we will surely find that about half the time the coin lands on the true side.

6.2. Random Sampling again

Four out of five chocolates in the box have cherry fillings, while the rest have caramel. Picking one at random, what should my credence be that it is cherry-filled? Everyone, including the staunchest opponents of POI, seem to agree on the answer 4/5. Now of course the chocolate I've chosen has many other features, for example this one is circular with a swirl on top. Noticing such features could hardly make a difference to my reasonable credence that it is cherry filled (unless of course I have some information regarding the relation between chocolate shapes and fillings). Often chocolate fillings do correlate with their shapes, but I haven't the faintest clue how they do in this case or any reason to suppose they correlate one way rather than another. But now for any probability function P ,

$$P(\text{cherry} \mid \text{circular}) = P(\text{cherry}) \leftrightarrow P(\text{cherry} \mid \text{circular}) - P(\text{cherry} \mid \text{-circular})$$

and my credence will meet the RHS condition only if it is also such that

$$P(\text{all and only the circular ones are cherry}) = P(\text{all and only the circular ones are caramel})$$

So I can have credence 4/5 the chocolate I've chosen is cherry-filled only if I give equal credence to these two hypotheses concerning shape-filling correlations. Should I give these equal credence? It is not as though I have any information concerning the objective chances of circular chocolates being filled a certain way. I don't know that the chocolate designers flipped a fair coin to decide the matter. It's just that I have no more reason to suppose that they would put one filling in one shape than another. In treating the shape as irrelevant and having 4/5 credence that this randomly selected chocolate is cherry-filled it appears we are implicitly committed to POI.

In a way that will be familiar by now, a proponent of the Chance Grounding Thesis will insist that my credence that it is cherry rather than caramel that goes in the circular chocolates should be mushy. The further result is that while my credence that the chosen chocolate is cherry-filled should be 4/5 prior to viewing it, once I see its shape (whatever shape it happens to be) my credence that it is cherry-filled should dilate to become mushy. But this is just not the way we think about such matters. And it's not just about chocolates. We are constantly faced with the possibility of correlations between attributes about which we know nothing. If we were not implicitly committed to certain

credences in accordance with POI then our credences would be mushy all over the place in implausible ways.

6.3. Updating on Mushy Credence

There are two kinds of chemicals used in this lab, X and Y . We know nothing about how common either chemical is or how this test tube came to be filled. But we have no evidence either way bearing on whether it contains X or Y . As usual POI says that $C(\text{contains } X) = 1/2$, the Grounding thesis says your credence should be mushy, perhaps even $C(\text{contains } X) = [0, 1]$. But whatever opinion we begin with it needn't stay that way. We should change our opinion in the light of new evidence. Let's run a test that has a 90% accuracy rate (i.e. whether it contains X or Y , the test has a 90% chance of accurately reporting whether it is X or Y). When we get a result, how confident should we be that it is accurate? Obviously 90%. Of course we do have to be careful in such cases that we do not commit a base rate fallacy. If we knew that X was far more common than Y , say, then getting a test result of Y would be some evidence that the result is inaccurate. But we have no such base rate information in this case, nor any reason to suppose there is one base rate or another. Now for any probability function P ,

$$\begin{aligned}
 P(\text{accurate result} \mid \text{result} : X) &= P(\text{contains } X \mid \text{result} : X) \\
 &= P(\text{contains } X)P(\text{result} : X \mid \text{contains } X) / [P(\text{contains } X) \\
 P(\text{result} : X \mid \text{contains } X) + P(\neg \text{contains } X)P(\text{result} : X \mid \neg \text{contains } X)] \\
 &= P(\text{contains } X) \cdot .9 / [P(\text{contains } X) \cdot .9 + P(\neg \text{contains } X) \cdot .1] \\
 &= .9 \text{ iff } P(\text{contains } X) = 1/2
 \end{aligned}$$

(Of course the same result holds for $\text{result} : Y$.)

Upon receiving the test result I can have credence of .9 that it was accurate only if I give equal credence to the two hypotheses concerning the contents of the test tube. POI scores again!

What happens if I have functions in my representor assigning multiple values to $\text{contains } X$ in my representor and update each by conditionalizing on the test result? Well, as Susanna Rinard pointed out to me, if prior to evidence e my credence in p covers the entire interval $[0,1]$, then no matter what evidential bearing e has on p , once I conditionalize on e my credence will still cover the range $[0,1]$. Maximally mushy credences are immovable!²¹ This result is entirely unacceptable. To avoid it we must either abandon the rule of updating each function by conditionalization, or deny that it is ever reasonable for credences to be *that* mushy. Neither option is appealing without forfeiting some of the motivating grounds for the mushy credence approach

²¹ With the exception of the extremities 0 and 1, each function P may increase its value for p , i.e. $P_+(p) > P(p)$. But then for any P such that $P(p) = x \in (0, 1)$, there is a P' such that $P'(p) < x$ but by conditionalization $P'_+(p) = x$. So we still have the same spread of values $C_+(p) = [0, 1]$.

to symmetrical ignorance. In what follows I will suppose that the opponent of POI continues with the rule of conditionalization but recommends a credal state of, let's say, $C(\text{contains } X) = [.05, .95]$ (never mind how we might arrive at these values). By my calculations, upon conditionalizing on *result: X*, my new credence spread will be approximately $C_+(\text{contains } X) = [.32, .99]$. The spread of my credence shifts upward slightly and sharpens a bit, but is nothing like .9. Meanwhile my credence that the test result is *accurate* dilates from sharply .9 prior to seeing the test result, to [.32, .99] upon viewing the result regardless of whether the test result is X or Y. Neither judgment fits the way that we normally reason. We are constantly take evidence such as test results, expert opinion, or general testimony as a basis for fairly sharp credences on matters that initially we were ignorant about.

7. CONCLUSION

The motivating ideas behind POI cannot be ignored, nor can it be easily refuted. The mushy credence approach to ignorance is the most credible alternative but faces serious problems. Either it needs to be worked out in some new subtle way, or we owe Laplace an apology for deriding his principle.²²

REFERENCES

- Carnap, R. (1950). *Logical Foundations of Probability* (Chicago).
- Elga, A. (2000). 'Self-locating Belief and the Sleeping Beauty Problem', *Analysis*, 60: 143–7.
- (MS). 'Subjective Probabilities Should be Sharp'.
- Gillies, D. (2000). *Philosophical Theories of Probability* (London).
- Hacking, I. (1965). *The Logic of Statistical Inference* (Cambridge).
- (2000). *An Introduction to Probability and Inductive Logic* (Cambridge).
- Hájék, A. (2007). 'Interpretations of Probability', *Stanford Encyclopedia of Philosophy*.
- Jaynes, E. T. (1973). 'The Well Posed Problem', *Foundations of Physics*, 3: 477–93.
- Joyce, J. (2005). 'How Probabilities Reflect Evidence', *Philosophical Perspectives*, 19: 153–78.
- Kyburg, H. (1970). *Probability and Inductive Logic* (London).
- (1974). *The Logical Foundations of Statistical Inference* (Dordrecht).
- (1977). 'Randomness and the Right Reference Class', *Journal of Philosophy*, 74: 501–21.
- Laplace, P. S. (1814/1951). *A Philosophical Essay on Probabilities* (New York).
- Levi, I. (1974). 'On Indeterminate Probabilities', *Journal of Philosophy*, 71: 391–418.
- (1977). 'Direct Inference', *Journal of Philosophy*, 74: 5–29.
- Lewis, D. (1980). 'A Subjectivist's Guide to Objective Chance', in R. Carnap and R. Jeffrey (eds.), *Studies in Inductive Logic and Probability*, ii (Berkeley).

²² Those who would rather pursue the first option could start with Scott Sturgeon's contribution to this volume (Ch. 5).

- McGrew, T. (2003). 'Direct Inference and the Problem of Induction', in H. Kyburg and M. Thalos (eds.), *Probability Is the Very Guide of Life* (Chicago).
- North, J. (MS). 'An Empirical Approach to Symmetry and Probability'.
- Seidenfeld, T., and Wasserman, L. (1993). 'Dilation for Sets of Probabilities', *Annals of Statistics*, 21: 1139–54.
- Sober, E. (2003). 'An Empirical Critique of Two Versions of the Doomsday Argument—Gott's Line and Leslie's Wedge', *Synthese*, 135: 415–30.
- Strevens, M. (1998). 'Inferring Probabilities from Symmetries', *Noûs*, 32: 231–46.
- Sturgeon, S. (2008). 'Reason and the Grain of Belief', *Noûs*, 42: 139–65.
- Talbot, W. (1991). 'Two Principles of Bayesian Epistemology', *Philosophical Studies*, 62: 135–50.
- van Fraassen, B. (1984). 'Belief and the Will', *Journal of Philosophy*, 81: 235–56.
- (1989). *Laws and Symmetry* (Oxford).
- (1990). 'Figures in a Probability Landscape', in J. Dunn and A. Gupta (eds.), *Truth or Consequences* (Dordrecht).
- (2005). 'Conditionalizing on Violated Bell's Inequalities', *Analysis*, 65: 27–32.
- (2006). 'Vague Expectation Loss', *Philosophical Studies*, 127: 483–91.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities* (London).
- Williamson, T. (2000). *Knowledge and its Limits* (Oxford).
- (2008). 'Why Epistemology Can't Be Operationalized', in Q. Smith (ed.), *New Essays in Epistemology* (Oxford).
- White, R. (2005). 'Epistemic Permissiveness', *Philosophical Perspectives*, 19: 445–59.