

Principles of Indifference

Benjamin Eva*

Forthcoming in *The Journal of Philosophy*

Abstract

The principle of indifference (PI) states that in the absence of any relevant evidence, a rational agent will distribute their credence (or ‘degrees of belief’) equally amongst all the possible outcomes under consideration. Despite its intuitive plausibility, PI famously falls prey to paradox, and so is widely rejected as a principle of ideal rationality. Some authors have attempted to show that by conceiving of the epistemic states of agents in terms of *imprecise credences*, it is possible to overcome these paradoxes and thus to achieve a consistent rehabilitation of PI. In this article, I present an alternative rehabilitation of PI in terms of the epistemology of *comparative confidence judgements* of the form ‘I am more confident in the truth of p than I am in the truth q ’ or ‘I am equally confident in the truth of p and q ’. In particular, I consider two natural comparative reformulations of PI, and argue that while one of them prescribes the adoption of patently irrational epistemic states, the other (which is only available when we drop the standard but controversial ‘Opinionation’ assumption from the comparative confidence framework) provides a consistent formulation of PI that overcomes the fundamental limitations of all existing formulations.

*Department of Philosophy, University of Konstanz, 78464 Konstanz (Germany) –

<http://be0367.wixsite.com/benevaphilosophy> – benjamin.eva@uni-konstanz.de.

1 Introduction

Bayesian epistemology is often characterised by the core constraints that (1) the epistemic states of rational agents should be conceived of in terms of precisely quantified numerical ‘degrees of belief’ or ‘credences’, (2) a rational agent’s credences should always be probabilistic, and (3) upon learning new evidence, rational agents should update their credences by means of Bayesian conditionalisation.¹ But these norms alone are not sufficient to define a universally applicable epistemic strategy for rational agents. Most pertinently, (1) (2) and (3) tell us nothing about what initial credences agents should adopt at the very beginning of their credal lives, before they obtain any relevant evidence about the world.² The best known solution to this problem is given by the *principle of indifference*,

Principle of Indifference (PI): Let $X = \{x_1, x_2, \dots, x_n\}$ be a partition of the set W of possible worlds into n mutually exclusive and jointly exhaustive possibilities. In the absence of any relevant evidence pertaining to which cell of the partition is the true one, a rational agent should assign an equal initial credence of $\frac{1}{n}$ to each cell.³

Armed with PI, the Bayesian now has access to a complete recipe for rationality that instructs agents not only on how to revise their credences in the face of new evidence, but also on what credences they should adopt in the absence of all evidence. And of course, PI is an extremely

¹i.e. upon learning evidence E , an agent should assign any proposition X a posterior degree of belief $P(X|E) = \frac{P(X \wedge E)}{P(E)}$.

²Other than that those credences should be probabilistic.

³For simplicity, I restrict my analysis to the finite case throughout the article. Since the paradoxes in question already arise in the finite setting, I take this to be an admissible simplification for present purposes.

intuitive and plausible principle with a number of independent justifications.⁴

Sadly though, it turns out that PI leads immediately to paradox once one attempts to apply it across multiple partitions of possibility space at the same time. To illustrate, consider the following example (due to van Fraassen, 1989). A hypothetical factory manufactures qualitatively identical cubes. The only information you are given is that the length of a side of any cube is at most 2 feet. This implies that the length of a side is either between 0 and 1 ($L_{0,1}$), or between 1 and 2 ($L_{1,2}$). These possibilities define a partition $X_1 = \{L_{0,1}, L_{1,2}\}$, and PI requires you to distribute your credence equally across its cells. Specifically, your prior probabilistic credal state Pr should satisfy the equality $Pr(L_{0,1}) = 1/2 = Pr(L_{1,2})$.

But instead of asking about the length of a side, you could equally well ask about the area of a side. The fact that the length of any side is at most 2 feet entails that the area of any side is at most 4 square feet. So we can partition the possibility space into the four possibilities: $X_2 = \{A_{0,1}, A_{1,2}, A_{2,3}, A_{3,4}\}$. Again, the evidence gives you no reason to prefer any one of these possibilities to any other, so PI requires you to assign them all equal prior probability, $Pr(A_{0,1}) = Pr(A_{1,2}) = Pr(A_{2,3}) = Pr(A_{3,4}) = 1/4$.

Now note that the sentences $L_{0,1}$ and $A_{0,1}$ are *logically equivalent*. But PI has required us to assign them different probabilities (1/2 and 1/4, respectively). Following PI, we've assigned two contradictory credences to a single proposition.

This paradox, of which there are many versions, has puzzled philosophers for roughly a century. One popular prospective solution (advocated by Keynes (1921) and Kass and Wasserman (1996), amongst others) has been to argue that there are certain 'privileged' partitions over which rational agents should apply PI. For example, it might be argued that in the cube example, you should only apply PI to the finer partition X_2 , and then derive your credences over X_1 from the distribution

⁴PI has been justified, for example, by considerations concerning epistemic utility (Pettigrew, 2014), risk aversion (Williamson, 2007), evidential support (White, 2009), informativity (Jaynes, 1957), and the Principal Principle (Hawthorne *et al*, 2015).

over X_2 . However, the details of such a resolution have remained problematic, and the philosophical foundations of this approach have been extensively critiqued in the literature (see e.g. Von Mises (1951), Norton (2008), Shackel (2007)). Here, I will assume that any satisfactory rehabilitation of PI will require rational agents to treat all possible partitions of possibility space equally, i.e. it will legislate that they should adopt the same kinds of epistemic attitudes across all partitions for which they lack any relevant evidence.

Another influential approach to resolving the paradox is to abandon the idea that the epistemic states of rational agents should always be conceived of in terms of precise numerical credences. Authors including Joyce (2005, 2010), Weatherson (2007), Kaplan (1996), Walley (1991) and Dupin (2015) have argued that PI can be consistently reformulated in a philosophically satisfactory manner once we allow for the possibility that rational agents sometimes have *imprecise* credences. Although the details of these approaches vary significantly between authors, the general idea is that an agent without access to any evidence regarding which cell of a partition is the true one should assign every cell *maximally imprecise credence*. This preserves the basic idea that indifferent agents should treat every cell of the partition equally, but it avoids paradox by rejecting the assumption that this equal treatment should be represented by an assignment of equal point valued credences.

But despite their popularity, these imprecise reformulations of PI are also plagued with fundamental technical and philosophical problems. In a recent series of articles, Rinard (2013, 2014) shows that most extant variations of the imprecise approach to PI (i) end up requiring agents to be indifferent between propositions about which they clearly should not be indifferent, and (ii) typically render meaningful inductive learning impossible.⁵ She also shows that the most natural response to these problems (to prescribe non-maximally imprecise credences) is unable to help, since it leads to probabilistic incoherence. Although my aim here is not to present an evaluation of imprecise reformulations of PI, I take it that these results cast enough doubt on the imprecise approach to motivate the search for an alternative solution to the paradoxes of indifference.

⁵This limitation was already noted by e.g. Joyce (2010).

Like the proponents of the imprecise reformulation of PI, I will attempt to resolve the paradoxes of indifference in the context of a more coarse grained representation of the epistemic states of rational agents. But rather than replacing precise numerical credences with imprecise interval valued credences, I will work within the framework of *comparative confidence judgements*. Specifically, I will assume that an agent's epistemic state is characterised by judgements of the form 'I am more confident in the truth of p than I am in the truth of q ' or 'I am equally confident in the truth of p and q '. Numerous eminent figures from the history of probability such as Keynes (1921), de Finetti (1937, 1951), Koopman (1940) and Fine (1973) have argued that comparative confidence judgements are the most fundamental, intuitive and psychologically basic of all our epistemic attitudes, so I am in good company when I assume that a rational agent's epistemic state can at least some times be fully characterised in comparative terms.

Furthermore, by studying possible reformulations of PI in terms of comparative confidence judgements, we can hope to gain a new perspective on the scope and tenacity of the paradoxes of indifference. For, comparative confidence judgements are far more coarse grained and generally applicable epistemic attitudes than precise assignments of numerical credence. If it turned out that the paradoxes of indifference still arose for agents whose epistemic states are characterised entirely by comparative confidence judgements, then the problem would be far more severe than has hitherto been recognised. Conversely, if we are able to obtain a philosophically satisfactory comparative reformulation of PI, that would help us to diagnose the root cause of the paradoxes of indifference, and could also be viewed as *prima facie* evidence for the claim that the norms of ideal rationality should sometimes be articulated as constraints on comparative confidence judgements.

The structure of the paper is as follows. In section 2 I present a concise overview of some of the most important technical assumptions and synchronic rationality constraints from the literature on comparative confidence judgements. In section 3 I consider the most obvious reformulation of PI in terms of comparative confidence judgements and demonstrate that, despite being consistent, it provides agents with intuitively irrational advice. In section 4 I show how one can obtain an

alternative comparative formulation of PI by dropping the controversial Opinionation assumption for comparative confidence orderings, and argue that the resulting formulation provides consistent, philosophically principled and intuitively rational epistemic guidance to agents. In section 5 I show how agents that organise their initial epistemic states in line with my preferred comparative reformulation of PI can undergo genuine inductive learning as they acquire new evidence. Section 6 concludes.

2 Comparative Confidence Orders

Let's begin with some formal preliminaries. Firstly, I assume that agents make comparative confidence judgements about 'propositions' drawn from some fixed Boolean algebra \mathfrak{B} of equivalence classes of logically equivalent sentences of a fixed language L . An agent A can make two kinds of comparative confidence judgement about propositions in \mathfrak{B} . They can be strictly more confident in the truth of p than they are in the truth of q . I denote this kind of judgement with the notation $[p \succ q]$. Alternatively, A can be equally confident in the truth of p and q . I denote this second kind of judgement with the notation $[p \sim q]$. The set of all A 's comparative confidence judgements defines a *confidence ordering*, \succsim , over some subset of the propositions in \mathfrak{B} , defined by $p \succ q$ if and only if $[p \succ q]$, and $p \sim q$ if and only $[p \sim q]$. I write $p \succsim q$ to indicate the disjunction ' $p \succ q$ or $p \sim q$ '. Most authors make the following basic assumption about \succsim .

Opinionation: For any $p, q \in \mathfrak{B}$, A makes exactly one of the judgements $[p \succ q]$, $[q \succ p]$, $[p \sim q]$,
i.e. one of $p \succ q$, $q \succ p$ and $p \sim q$ is true.

Opinionation implies that \succsim is a 'total ordering' of \mathfrak{B} . Intuitively, it means that there are 'no gaps' in A 's confidence judgements, i.e. A makes a comparative confidence judgement about every pair of propositions in \mathfrak{B} . This assumption, though controversial, is standard in the extant literature on comparative confidence orderings. It will play a crucial role in assessing prospective comparative reformulations of PI. It is also typically assumed that \succ satisfies the following conditions.

Irreflexivity of \succ : For every $p \in \mathfrak{B}$, A does not make the judgement $[p \succ p]$, i.e. $p \not\succeq p$.

Transitivity of \succ : For every $p, q, r \in \mathfrak{B}$, if A makes the judgements $[p \succ q]$ and $[q \succ r]$, then A makes the judgement $[p \succ r]$, i.e. if $p \succ q$ and $q \succ r$, then $p \succ r$.

Finally, it is standard to assume that \sim is an equivalence relation, i.e.

Reflexivity of \sim : For every $p \in \mathfrak{B}$, A makes the judgement $[p \sim p]$, i.e. $p \sim p$.

Transitivity of \sim : For every $p, q, r \in \mathfrak{B}$, if A makes the judgements $[p \sim q]$ and $[q \sim r]$, then A makes the judgement $[p \sim r]$, i.e. if $p \sim q$ and $q \sim r$, then $p \sim r$.

Symmetry of \sim : For every $p, q \in \mathfrak{B}$, if A makes the judgements $[p \sim q]$, then A makes the judgement $[q \sim p]$, i.e. if $p \sim q$, then $q \sim p$.

When all of these assumptions are satisfied, the ordering \succsim is called a ‘total preorder’ over \mathfrak{B} .

In the literature, it is almost universally assumed that the comparative confidence judgements of rational agents will always define a total preorder over \mathfrak{B} . However, there is still relatively little consensus regarding which additional properties are characteristic of rationally permissible confidence orderings. The most important additional constraints are normally phrased in terms of the representability of \succsim by numerical functions. Given a comparative confidence ordering \succsim over \mathfrak{B} and a function $\mu : \mathfrak{B} \rightarrow [0, 1]$, say that \succsim is ‘fully represented’ by μ if and only if for every $p, q \in \mathfrak{B}$, (i) $p \succ q \Leftrightarrow \mu(p) > \mu(q)$, and (ii) $p \sim q \Leftrightarrow \mu(p) = \mu(q)$. Icard (2016) gives a pragmatic ‘money-pump’ style argument in favour of the requirement that \succsim should always be fully representable by a probability function, while Fitelson and McCarthy (unpublished) provide an epistemic utility theoretic justification for the requirement that \succsim should always be fully representable by a Dempster Schafer belief function (they also show that no such justification can be given for the probabilistic representability requirement). Here, I remain agnostic about which (if any) of these representability conditions we should accept, and assume only the following almost universally accepted monotonicity conditions.

$$(A1) \top \succ \perp.$$

(A2) For any $p, q \in \mathfrak{B}$, if $p \vdash q$ then $q \succsim p$.

A1 requires that rational agents always be strictly more confident in the tautology than they are in the contradiction, and A2 is a general monotonicity requirement, which stipulates that agents should never be strictly more confident in p than they are in the logical consequences of p . As well as being intuitively compelling, these rationality constraints have been given a range of pragmatic justifications (see e.g. Fishburn (1986), Halpern (2003)). Here, we will assume only that rational agents always have comparative confidence orderings that (i) define total preorders over \mathfrak{B} , and (ii) satisfy A1 and A2. We are now ready to translate PI to this very general setting.

3 Comparative Indifference: Take 1

We've seen that the standard probabilistic version of PI leads straight to paradox when we attempt to apply it simultaneously across multiple partitions of possibility space. But now suppose that we are trying to advise a rational agent whose epistemic state is characterised purely by a comparative confidence ordering of the type described in the previous section. What comparative confidence judgements should such an agent make when they do not have access to any relevant evidence regarding the possibilities under consideration? The most natural suggestion is the following, which is a direct translation of PI to the comparative setting.

Comparative Principle of Indifference 1 (CPI₁): Let $X = \{x_1, x_2, \dots, x_n\}$ be a partition of the set W of possible worlds into n mutually exclusive and jointly exhaustive possibilities. In the absence of any relevant evidence pertaining to which cell of the partition is the true one, a rational agent should make the comparative confidence judgements $[x_i \sim x_j]$, for all $i, j \leq n$.

Clearly, CPI₁ appears to capture the essential philosophical intuition behind PI, i.e. that you shouldn't prefer one outcome to another unless you have evidence telling you to do so. It prescribes

that a rational agent without access to any evidence pertaining to which cell of the partition is the true one should be equally confident in every cell.

The fact that a comparative reformulation like CPI_1 provides perfectly consistent guidance to rational agents was first noted by Norton (2008, 2010).⁶ To illustrate, recall again van Fraassen’s cube example. In this case there are two relevant partitions over which the agent should apply CPI_1 . These are $X_1 = \{L_{0,1}, L_{1,2}\}$ (the partition in terms of length) and $X_2 = \{A_{0,1}, A_{1,2}, A_{2,3}, A_{3,4}\}$ (the partition in terms of area). CPI_1 recommends that the indifferent agent should be equally confident in the two hypotheses regarding length and equally confident in the four hypotheses regarding area. But since $A_{0,1}$ and $L_{0,1}$ are logically equivalent, this immediately implies (using the fact that \sim is an equivalence relation) that the agent is equally confident in all of the propositions being considered, i.e. their ordering satisfies the following relations:

$$L_{0,1} \sim L_{1,2} \sim A_{0,1} \sim A_{1,2} \sim A_{2,3} \sim A_{3,4}$$

There is no ambiguity or inconsistency in the epistemic advice given by CPI_1 . The principle tells the agent exactly which epistemic attitudes to adopt, and at first blush, these attitudes look plausible. So one might think that the paradox of indifference simply vanishes in the comparative setting. As Norton puts it,

[I]f we forgo the idea that belief states must be probability distributions, then there is a unique, well defined, ignorance state. It assigns the same, unique, ignorance degree of belief to every contingent outcome and each of their contingent disjunctive parts in the outcome space. (Norton, 2008: 46)

However, I contend that something has gone seriously wrong in the comparative reformulation of PI described above. To see this, note first that the sentence $L_{1,2}$ is logically equivalent to the

⁶Norton actually doesn’t formalise the principle in terms of comparative confidence judgements, but rather in terms of ‘degrees of inductive support’, which are not assumed to be probabilistic. It is easy to see that the reformulation used here is formally equivalent to Norton’s version of PI.

disjunction $A_{1,2} \vee A_{2,3} \vee A_{3,4}$.⁷ And CPI_1 has recommended that rational agents without any evidence regarding the area of the sides of the cube should make the judgement $L_{1,2} \sim A_{1,2}$, i.e.

$$A_{1,2} \sim (A_{1,2} \vee A_{2,3} \vee A_{3,4})$$

Informally, this means that CPI_1 is advising the agent to be equally confident in the claims ‘the area of a side of the cube is between 1 and 2 square feet’ and ‘the area of a side of the cube is between 1 and 4 square feet’ when they have no absolute no relevant evidence to go on. But this seems deeply irrational. To make this intuitive irrationality precise, consider the following plausible principle for connecting comparative confidence judgements to betting behaviour.

Indifferent Betting Principle (IBP): Given two proposition p and q , and two gambles of the form

Gamble 1: o_{good} if p is true, and o_{bad} if p is false

Gamble 2: o_{good} if q is true, and o_{bad} if q is false

if a rational agent A makes the judgement $p \sim q$, then they should not be willing to pay any positive price to exchange gamble 1 for gamble 2 (or vice-versa).

I won’t do much here to defend IBP, which I take to be a prima-facie compelling principle.⁸ Intuitively, it simply says that if you are equally confident in p and q , then you shouldn’t pay anything to exchange a gamble on p for a gamble on q with the same possible payoffs, which strikes me as an obvious and fundamental edict of practical rationality.

Now, imagine that we tell an agent about the cube factory and offer them the choice between the following two options:

⁷A square has area between 1 and 4 square feet if and only if it’s sides have length between 1 and 2 feet.

⁸A structurally similar assumption regarding the relationship between comparative confidence judgements and betting behaviour is made by Icard (2016). Note that if the comparative confidence ordering is derived from an agent’s probabilistic credences, then IBP is a logical consequence of maximising expected utility.

(1) Take the following gamble for free:

Gamble 1: Win \$1,000,000 if the area of a side is between 1 and 2 square feet, and win \$0 otherwise.

(2) Pay *any* positive price $\alpha > 0$ to buy the following gamble:

Gamble 2: Win \$1,000,000 if the area of a side is between 1 and 4 square feet, and win \$0 otherwise.

Next, suppose that the agent follows the advice given by CPI_1 and conforms to the betting principle IBP. Then they will always take option (1). Intuitively, this is deeply irrational. The agent is not willing to pay any price, *however small* in order to greatly increase the size of the region of evidentially possible worlds in which they would win \$1,000,000, in spite of the fact that they have absolutely no evidence that speaks against the plausibility of the worlds that lie in the region they're neglecting. If that's what the principle of indifference recommends, then it should be rejected as a principle of ideal rationality.

It should be noted that Norton (2008) actually imposes the condition that indifferent agents should be equally confident in the propositions 'the area is between 1 and 2 square feet' and 'the area is between 1 and 4 square feet' as a *desideratum* for any prospective formulation of epistemic indifference,⁹ and takes the fact that the comparative reformulation CPI_1 satisfies this desideratum to count in its favour. The preceding analysis demonstrates, contra Norton, that any formulation of PI that has this property should be rejected on the grounds that it prescribes manifestly irrational epistemic attitudes.

We've seen that the comparative reformulation of PI given by CPI_1 , unlike its probabilistic counterpart, gives definite and consistent advice when applied simultaneously across multiple different partitions of possibility space. However, we've also seen that despite being consistent, this advice is bad. Given a very weak and intuitively compelling principle linking comparative confi-

⁹This is entailed by his requirement that indifference should be 'invariant across disjunctive coarsenings of the outcome space'.

dence judgements to betting behaviour, an agent that adheres to CPI_1 will behave in strange and patently irrational ways.

4 Comparative Indifference: Take 2

Happily, there is another way to capture the basic philosophical intuition behind PI in a comparative setting. To see this, let's turn our attention briefly to the Opinionation assumption that I presented at the beginning of Section 2. According to this assumption, for any two propositions p and q , a rational agent should always either (i) be more confident in p than they are in q , (ii) be equally confident in p and q , or (iii) be more confident in q than they are in p . Of course, if an agent's comparative confidence judgements are derived from an assignment of precise probabilistic credences, then their confidence ordering will always satisfy Opinionation. But recall that we are working in a generalised setting in which it is not assumed that agents have well defined credences. We are assuming that it is possible for the content of a rational agent's epistemic state to be exhausted by their comparative confidence judgements. Once we make this assumption, it is not clear why we should assume that rationally permissible confidence orderings always satisfy Opinionation. Indeed, dissatisfaction with the Opinionation assumption has been expressed by many influential authors. Thus, we read, for example (my emphases)

It seems to me that at present it is not yet clear whether the concept of degree of confirmation can be defined satisfactorily as a quantitative concept... Perhaps it is preferable to define it as a merely topological concept, i.e. by defining only the relations: 'S1 has the same or a higher degree of confirmation than S2, respectively' but in such a way that *most of the pairs of sentences will be incomparable*. (Carnap 1936: 427)

I maintain...that there are some pairs of probabilities between the members of which *no comparison of magnitude is possible*. (Keynes 1921: 36)

The implausibility of Opinionation as a general principle of ideal rationality is often cited as

a motivation for the replacement of standard Bayesianism with its imprecise counterpart (see e.g. Kaplan, 1996). Here, I won't go into details regarding the usual reasons for rejecting Opinionation (see e.g. Keynes 1921, Forrest 1989 for characteristic discussions). Rather, I will show that, by dropping the Opinionation assumption, we put ourselves in a position to obtain an alternative comparative reformulation of PI that, unlike its predecessor, does not give systematically bad advice. I take this to constitute a strong argument in favour of rejecting Opinionation, *in addition* to those that already exist in the literature.

To make things concrete, we need to slightly amend the formal framework described in section 2. Recall that the collection of an agent's comparative confidence judgements define a confidence ordering \succsim over the algebra \mathfrak{B} of possible outcomes. Until now we've assumed that \succsim is a total preorder on \mathfrak{B} . By dropping only the Opinionation assumption, we allow for the possibility that \succsim can be a *partial preorder* over \mathfrak{B} . This means that there can be $p, q \in \mathfrak{B}$ such that neither $p \succsim q$ nor $q \succsim p$ are true in the ordering. In such a case, we write $p \odot q$ and say that p and q are 'incomparable' in the ordering. Conceptually, it is important to note that incomparability is not another kind of comparative confidence judgement, but rather *the absence of any comparative confidence judgement whatsoever*. If p and q are incomparable in an agent's confidence ordering, then the agent has simply refrained from making any judgement regarding their comparative plausibility. It is also important to note that dropping Opinionation does nothing to 'infect' the other basic assumption we made about the structure of the ordering. It is still very plausible to assume that the 'strictly more confident than' is transitive and a-symmetric, that the 'equally confident in' relation is an equivalence relation, and that rational agents (i) should be strictly more confident in the tautology than they are in the contradiction, and (ii) should always be at least as confident in the logical consequences of p as they are in p . Since we've simply defined incomparability as the complement of all other possible ordering relations, it trivially follows that \odot is symmetric, irreflexive and typically non-transitive.

It is also important to ask whether there are any additional synchronic rationality constraints

that should be imposed on comparative confidence orderings in the absence of the Opinionation assumption. This question has not yet been systematically explored in the philosophical literature (although see Forrest (1989) for an illuminating discussion). One formal condition that has been tentatively forwarded in the technical literature (see e.g. Harrison-Taylor *et al* (2016), Forrest (1989), Alon and Lehrer (2014)) is the following:

(A3) there should exist a non-empty set \mathcal{P} of additive probability measures on \mathfrak{B} such that for

any $p, q \in \mathfrak{B}$:

$$p \succsim q \Leftrightarrow (\forall P \in \mathcal{P})(P(p) \geq P(q))$$

$$p \odot q \Leftrightarrow (\exists P, P' \in \mathcal{P})(P(p) > P(q) \wedge (P'(p) < P'(q)))$$

When \succsim satisfies A3, I say that it is ‘fully represented’ by any set \mathcal{P} that makes A3 true. It is important to note that assuming A3 is certainly *not* equivalent to assuming that an agent’s epistemic state is characterised by any particular set of probability functions (a la Joyce (2005, 2010) and Weatherson (2007)). Generally, if a confidence ordering satisfies A3, there will be infinitely many different sets of probability functions that could be used to represent it, and no particular set has any special claim to representing the agent’s epistemic state.

At any rate, A3 is a much stronger condition than is needed for present purposes. In what follows, I assume only the extremely weak and plausible condition

$$(A4) \text{ If } p \succsim q \text{ then } \neg q \succsim \neg p$$

A4 requires only that if I’m at least as confident in the truth of p as I am in the truth of q , then I should be at least as confident in the falsity of q as I am in the falsity of p . Again, I won’t say much to motivate this condition, which I take to be a prima-facie compelling principle of ideal rationality. However, it is worth noting that in the Opinionated setting, A4 is implied by probabilistic representability (but not by Dempster-Schafer belief function representability), and in the non-Opinionated setting, it is implied by A3. In what follows, I will mainly use examples of orderings that satisfy A3, since seeing the orderings represented by sets of probability distributions

is likely to aid the reader’s intuitive understanding. But I stress that this is *not* indicative of any substantive philosophical commitments on my part. A4 is the only substantive synchronic rationality constraint that plays a role in the subsequent analysis, and it does not require that the confidence orderings of rational agents be representable by sets of probability distributions.

We are now ready to see how this generalised comparative framework allows us to obtain another plausible comparative reformulation of PI.

Comparative Principle of Indifference 2 (CPI₂): Let $X = \{x_1, x_2, \dots, x_n\}$ be a partition of the set W of possible worlds into n mutually exclusive and jointly exhaustive possibilities. In the absence of any relevant evidence pertaining to which cell of the partition is the true one, a rational agent should not make any comparative confidence judgements regarding the pairs (x_i, x_j) for $i, j \leq n$. Equivalently, for all $i, j \leq n$, $x_i \odot x_j$ should hold in the agent’s confidence ordering.

Like CPI₁, CPI₂ captures the basic intuition that an agent without access to any relevant evidence should treat every cell of the partition equally. In particular, they should not regard any cell as more plausible than any other. More than that, they should simply refrain from making any comparative confidence judgements *whatsoever* regarding the cells of the partition. Intuitively, CPI₂ tells agents not to make any comparative confidence judgments (including judgments of equal plausibility) until they encounter evidence supporting those judgments.

However, unlike CPI₁, CPI₂ does not force an agent to adopt the kind of radically non-additive and patently irrational epistemic states prescribed by CPI₁. To illustrate, consider again van Fraassen’s cube example. CPI₂ requires that the indifferent agent should refrain from making any comparative confidence judgements regarding the cells in the partitions $X_1 = \{L_{0,1}, L_{1,2}\}$ and $X_2 = \{A_{0,1}, A_{1,2}, A_{2,3}, A_{3,4}\}$, respectively. This amounts to requiring that their confidence ordering should satisfy the relations


$$(i) L_{0,1} \odot L_{1,2}$$

$$(ii) A_{0,1} \odot A_{1,2}, A_{0,1} \odot A_{2,3}, A_{0,1} \odot A_{3,4}, A_{1,2} \odot A_{2,3}, A_{1,2} \odot A_{3,4}, A_{2,3} \odot A_{3,4}$$

Note that (i) and (ii) do *not* imply that $A_{1,2} \odot L_{1,2}$ even though we have $L_{1,2} \odot L_{0,1} \odot A_{1,2}$. For, unlike the ‘equally confident’ relation (\sim), the incomparability relation (\odot) is typically not transitive. This is crucial, since just as I argued that it would be deeply implausible for an agent to be equally confident in the propositions ‘the area is between 1 and 2’ and ‘the area is between 1 and 4’, it would seem to be equally implausible for an agent to simply refrain from making any comparative confidence judgement about those propositions. Surely a rational agent should be *more* confident in the latter proposition than they are in the former, for the reasons described in the previous section. It’s easy to see that this requirement is perfectly consistent with CPI_2 . For example, consider the ordering \succsim that corresponds (in the sense defined by $A3$) to the set of probability distributions $\mathcal{P} = \{P_1, P_2, P_3, P_4\}$, where

$$\begin{aligned} P_1(A_{0,1}) &= 0.7, P_1(A_{1,2}) = 0.1, P_1(A_{2,3}) = 0.1, P_1(A_{3,4}) = 0.1, \\ P_2(A_{0,1}) &= 0.1, P_2(A_{1,2}) = 0.7, P_2(A_{2,3}) = 0.1, P_2(A_{3,4}) = 0.1, \\ P_3(A_{0,1}) &= 0.1, P_3(A_{1,2}) = 0.1, P_3(A_{2,3}) = 0.7, P_3(A_{3,4}) = 0.1, \\ P_4(A_{0,1}) &= 0.1, P_4(A_{1,2}) = 0.1, P_4(A_{2,3}) = 0.1, P_4(A_{3,4}) = 0.7 \end{aligned}$$

It is easy to see that so defined, \mathcal{P} defines the confidence ordering \succsim depicted below.



The first thing to note here is that this ordering straightforwardly satisfies CPI_2 for the cube example. None of the propositions in the area partition $X_2 = \{A_{0,1}, A_{1,2}, A_{2,3}, A_{3,4}\}$ is comparable to any other, as CPI_2 demands, and similarly for the length partition $X_1 = \{A_{0,1}, \neg A_{0,1}\}$ (where $A_{0,1}$ and $\neg A_{0,1}$ are equivalent to $L_{0,1}$ and $L_{1,2}$, respectively). Furthermore, the ordering ranks the proposition $A_{1,4}$ strictly above $A_{1,2}$, $A_{2,3}$ and $A_{3,4}$, which is exactly the desirable property that was violated when we followed the advice of CPI_1 . Agents can follow the advice of CPI_2 while making the common sense judgement that the proposition ‘the area is between 1 and 4 square feet’ is strictly more plausible than the proposition ‘the area is between 1 and 2 square feet’. So as well as giving us consistent advice, it seems that CPI_2 also gives us *epistemically reasonable* advice, unlike its predecessor, CPI_1 .

At this stage, the reader might notice that the confidence ordering depicted above has a familiar structure. Specifically, it has the structure of the 16 Boolean algebra that depicts the classical logical structure of a language with four mutually exclusive and jointly exhaustive atomic propositions.¹⁰ Thus, we have an example of a confidence ordering that (i) satisfies CPI_2 , (ii) is fully representable by a set of probability functions, and therefore satisfies A4, (iii) doesn’t encode any patently irrational comparative confidence judgements, and (iv) is actually isomorphic to the logical structure of the language being used to describe the possibility space. In fact, it turns out that this is an instance of a general phenomena, as the following theorem shows.

Theorem: Let \mathfrak{B} be the Boolean algebra encoding the logical structure of a language L . If an agent does not have access to any evidence pertaining to any non-trivial partitions of \mathfrak{B} , then the only way for them to satisfy both CPI_2 and A4 is for their confidence ordering over \mathfrak{B} to be isomorphic to \mathfrak{B} itself.¹¹

Proof: It follows immediately from its definition that \mathfrak{B} is always a partial pre-order over its

¹⁰Standardly referred to as the ‘Lindenbaum algebra’ of the language.

¹¹By a ‘non-trivial partition’, I mean a partition that consists only of contingent propositions.

own underlying set of propositions that satisfies $A1, A2$ and $A4$. To show that \mathfrak{B} satisfies CPI_2 , let $\{p_1, \dots, p_n\}$ be a partition of \mathfrak{B} . By definition, none of the p_i 's entail one another (they're mutually exclusive), so they are all incomparable in \mathfrak{B} 's ordering, as desired.

Conversely, let \succsim be an partial pre-order over \mathfrak{B} 's underlying set that satisfies $A1, A2, A4$ and CPI_2 . In order to show that \succsim is isomorphic to \mathfrak{B} , it suffices to show that $q \succsim p$ if and only if $p \vdash q$. By $A2$, $p \vdash q$ implies $q \succsim p$. Next, suppose for contradiction that $q \succsim p$ and $p \not\vdash q$. Since $p \not\vdash q$, $p \wedge \neg q$ is a contingent possibility and $\{p \wedge q, p \wedge \neg q, \neg p\}$ is a non-trivial partition. Since \succsim satisfies CPI_2 , it follows that $(p \wedge \neg q) \odot \neg p$. Furthermore, since \succsim satisfies $A4$ and $q \succsim p$, we have $\neg p \succsim \neg q \succsim (p \wedge \neg q)$, contradicting $(p \wedge \neg q) \odot \neg p$. ■

This theorem shows that, combined with the very weak synchronic rationality constraint $A4$, CPI_2 always prescribes a unique confidence ordering for agents at the very beginning of their credal lives, namely the ordering that simply encodes the deductive relationships between the propositions in the relevant language. This strengthens the argument for the superiority of CPI_2 as compared to CPI_1 . Whereas CPI_1 entails that agents in the cube example should always be equally confident in the propositions ‘The area is between 1 and 2 square feet’ and ‘The area is between 1 and 4 square feet’, CPI_2 (combined with $A4$) entails that agents should *always* be strictly more confident in the latter than they are in former. It is not just that CPI_2 permits agents to make such obviously rational judgements. It *requires* them to do so.

5 How to Raise a Superbaby

In the previous section, I argued that CPI_2 is the best available comparative reformulation of PI. However, one might be worried that the prospects of embedding CPI_2 in a plausible and non-trivial inductive logic are dim. Specifically, it's unclear how an agent whose initial epistemic state is characterised by a partial preorder over \mathfrak{B} should revise their epistemic attitudes as they gain new

evidence about the world. Here, I show that there is actually a very natural and straightforward way to embed CPI_2 in a general inductive logic that generalises Bayesian conditionalisation to the comparative setting.

To begin, consider the question ‘how should an agent revise their comparative confidence judgements over time as they gather new evidence about the world?’. More precisely, suppose that an agent begins with a confidence ordering \succsim and subsequently learns some evidential proposition e with certainty. They then need to adopt a new ordering that satisfies the relation $e \sim \top$ (reflecting their new found certainty in e) and takes into account all the evidential ramifications of their learning experience. How should they do this? One natural option is to take a cue from Bayesian conditionalisation, i.e. to study how Bayesian agents revise their comparative confidence judgements when they conditionalise on e . Given a probability distribution P , let \succsim_P be the confidence ordering defined by $q \succ_P p$ if and only if $P(q) > P(p)$ and $p \sim_P q$ if and only if $P(p) = P(q)$. By definition, P fully represents \succsim_P , and we can think of \succsim_P as encoding the comparative confidence judgements of a Bayesian agent whose credal state is given by the probability function P . Now, we can ask ‘what is the relationship between \succsim_P and $\succsim_{P(-|e)}$ ’, where $P(-|e)$ is the probability function obtained by conditionalising P on e . This question has a simple answer:

$$\begin{aligned}
q \succ_{P(-|e)} p &\Leftrightarrow P(q|e) > P(p|e) \\
&\Leftrightarrow P(q|e)p(e) > P(p|e)p(e) \\
&\Leftrightarrow P(e \wedge q) > P(e \wedge p) \\
&\Leftrightarrow e \wedge q \succ_P e \wedge p
\end{aligned}$$

Thus, if we let \succsim_e denote the ordering that results from revising the initial ordering \succsim after learning e , a Bayesian agent will always revise their confidence orderings according to the rule

$$\text{(CC:)} \quad q \succ_e p \Leftrightarrow e \wedge q \succ e \wedge p, \text{ and } q \sim_e p \Leftrightarrow e \wedge q \sim e \wedge p$$

Where ‘CC’ stands for ‘comparative conditionalisation’. In a companion paper (Eva, unpub-

lished), I consider the question of whether there is anything normatively special about CC from the perspective of the epistemology of comparative confidence judgements. It turns out that in contexts where Opinionation is assumed, CC can be justified both in terms of expected accuracy minimisation (utilising Fitelson and Mccarthy’s (unpublished) epistemic utility theoretic framework for comparative confidence orderings), and also in terms of a norm of epistemic conservativity. Of course, one of the central contentions of this article is that the Opinionation assumption is incompatible with a proper treatment of indifference, so we need to generalise the definition of CC, in which case these normative justifications may no longer apply. The obvious generalisation is the following.

$$\text{(CC*): } q \succ_e p \Leftrightarrow e \wedge q \succ e \wedge p, q \sim_e p \Leftrightarrow e \wedge q \sim e \wedge p \text{ and } q \odot_e p \Leftrightarrow e \wedge q \odot e \wedge p$$

The problem of providing normative justifications for this generalisation of CC to the non-opinionated setting is an important one that I leave aside for now. For present purposes, I am happy to simply posit CC* as the natural generalisation of a revision rule with strong normative justifications in the presence of the Opinionation assumption. It is also worth noting that whenever a non-opinionated ordering \succsim is fully represented by a set \mathcal{P} of probability functions, \succsim_e will always be fully represented by the conditionalized set $\{P(-|e) | P \in \mathcal{P}\}$.

To illustrate how CC* allows agents to move from the kind of indifferent epistemic states advocated by CPI₂ to more opinionated epistemic states as they gather new evidence, consider the following example (due to Keynes, 1921). We are interested in the nationality of a stranger, let’s call her ‘Sue’. We are told that Sue is definitely from one of Ireland, France or Great Britain, but we have no evidence regarding which of these is the true country. Thus, we are indifferent over both the maximally fine partition $\{F, GB, I\}$ and the disjunctively coarsened partition $\{F, GB \vee I\}$. CPI₂ requires that we should not make any comparative confidence judgements regarding the cells of either of these partitions. Now suppose that we also have some prior information regarding the popularity of rugby in France, Great Britain and Ireland. In particular, we know that rugby is very popular in Ireland, more so than in either England or France. Then, despite being indifferent about


Sue’s nationality, it is natural to expect that we should judge the proposition ‘Sue is Irish and likes rugby’ to be strictly more plausible than the propositions ‘Sue is French and likes rugby’ and ‘Sue is British and likes rugby’. A simple example of an ordering that satisfies all of these conditions is the ordering that is fully represented by the set $\mathcal{P} = \{P_1, P_2, P_3\}$, where

$$0.1 = P_1(F \wedge R) = P_1(GB \wedge R) = P_1(I \wedge \neg R), \quad 0.15 = P_1(GB \wedge \neg R) = P_1(I \wedge R), \quad 0.4 = P_1(F \wedge \neg R)$$

$$0.1 = P_2(F \wedge R) = P_2(GB \wedge R) = P_2(I \wedge \neg R), \quad 0.15 = P_2(F \wedge \neg R) = P_2(I \wedge R), \quad 0.4 = P_2(GB \wedge \neg R)$$

$$0.1 = P_3(I \wedge \neg R) = P_3(F \wedge R) = P_3(GB \wedge R), \quad 0.15 = P_3(F \wedge \neg R) = P_3(GB \wedge \neg R), \quad 0.4 = P_3(I \wedge R)$$

It is easy to see that relative to this ordering, $F \odot (GB \vee I)$ and $F \odot GB, F \odot I, GB \odot I$ all hold, as required by CPI_2 . However, the ordering also satisfies $I \wedge R \succ F \wedge R \sim GB \wedge R$. Now imagine that you learn that Sue does in fact like Rugby. Revising by CC^* leads to the posterior ordering \succsim_R represented below.



As expected, this ordering has more comparative structure than the prior ordering. I is now ranked above F and GB , representing the fact that we’ve updated on inductive evidence that supports the hypothesis that Sue is Irish.

Although the update procedure described here is so simple as to seem almost trivial, it is no small point in favour of CPI_2 that it can be straightforwardly integrated into a plausible inductive logic that allows for genuine inductive learning in all the cases that one would expect. For, one of the main criticisms of the representation of epistemic indifference as maximally imprecise credence

is that it typically renders meaningful inductive learning impossible since the prior credence of $[0, 1]$ is dogmatic, i.e. no matter how much inductive evidence one conditionalizes on, they will retain their prior imprecise credence of $[0, 1]$ in the hypothesis under consideration (see e.g. Rinard (2013), Joyce (2010)). If indifference is maximally imprecise credence, then agents are doomed to remain indifferent forever, regardless of how much evidence they obtain. This problem is completely avoided in the comparative treatment of indifference advocated here, which allows indifferent agents to obtain more structured epistemic states as they gather evidence about the world.

At this stage, it is worth pausing to preempt a possible reply that could be forwarded by advocates of the imprecise approach to indifference.¹² To illustrate the reply, consider the following example, due to Rinard (2013). You are presented with an urn. You know a-priori that the urn either contains only green marbles (H_1), or contains marbles of multiple colours, only $\frac{1}{10}$ th of which are green (H_2). Assuming you have no further evidence, you are indifferent between H_1 and H_2 . Suppose that you then go on to draw a random marble from the urn, which turns out to be green. This should increase your confidence in H_1 , as should each subsequent drawing of a green marble (up until the point at which you draw a non-green marble). The problem with the imprecise approach to indifference is that if we represent your prior indifference regarding H_1 by assigning H_1 the imprecise prior credence $[0, 1]$ and then conditionalise on the proposition $E =$ ‘a green marble is drawn’ to model your response to the subsequent evidence, your credence in H_1 will remain fixed at $[0, 1]$.¹³ This is implausible and unintuitive. Drawing green marbles should increase your confidence in H_1 (at least in comparison to H_2). The comparative approach to indifference defended here encounters no such problem. In order to model your indifference between H_1 and H_2 , we assume that H_1 and H_2 are incomparable in your prior confidence ordering. Next, we note that the conjunction $H_1 \wedge E$ is intuitively more plausible than $H_2 \wedge E$, i.e. that your prior confidence ordering should encode the judgement $H_1 \wedge E \succ H_2 \wedge E$. Since there is no reason to regard either one of H_1 or H_2 as

¹²Many thanks to an anonymous referee for raising this important objection.

¹³Assuming that all the credence functions in your imprecise ‘representor’ set respect the stipulated constraints that

$P(E|H_1) = 1$ and $P(E|H_2) = \frac{1}{10}$.

more plausible than the other, and E is clearly more plausible under the supposition that H_1 is true than it is under the supposition that H_2 is true, I take this to be a compelling assumption. We can then apply CC* to model the learning of E , and we get the desired result that H_1 is above H_2 in your posterior confidence ordering. Whereas the imprecise approach implausibly represents your credence in H_1 (and H_2) as being fixed across the learning experience, the comparative model has the simple virtue that it models you as becoming more confident in H_1 (in comparison to H_2) when you draw a green marble.

Here, the advocate of the imprecise approach will likely be tempted to reply that the comparative model only solves the inductive learning problem by helping itself to extra structure in the agent's prior epistemic state – structure that the imprecise approach does not assume. In particular, the comparative model assumes that your prior epistemic state encodes the judgement $H_1 \wedge E \succ H_2 \wedge E$. This judgement is not assumed by the imprecise representation, and the advocate of the imprecise approach may well consider its inclusion in the comparative model a kind of 'cheating'. More precisely, they may contend that if the imprecise representation were to be augmented with this assumption, then it too would be able to solve the inductive learning problem. So the purported progress made by the comparative model comes from assuming more structure in the agent's prior epistemic state, not from changing the way in which that state and its evolution are conceptualised.

Despite its initial plausibility, this response is ultimately unsuccessful. To see this, let's challenge the advocate of the imprecise approach to make good on their promise that adding the assumption $H_1 \wedge E \succ H_2 \wedge E$ as a constraint on the agent's prior epistemic state is sufficient to defuse their problem with inductive learning. The natural way to interpret this constraint in the imprecise model is as requiring that every probability function in the agent's 'representor' satisfies the constraint $P(H_1 \wedge E) > P(H_2 \wedge E)$.¹⁴ As it turns out, this assumption is incompatible with the assumption that the agent's prior credence in H_1 is the full closed interval $[0, 1]$, since it implies that every function

¹⁴Where an agent's 'representor' is the set of probability functions that is taken to represent their epistemic state on the imprecise model (see e.g. Joyce (2010)).

in the representor satisfies $P(H_1) > \frac{1}{10}P(H_2) \geq 0$. However, one can accommodate the assumption by slightly amending the imprecise approach to indifference so that it legislates a prior imprecise credence of $(0, 1)$ (as opposed to $[0, 1]$) in H_1 .¹⁵ But as it turns out, doing so does absolutely nothing to solve the imprecise model’s problem with inductive learning. Even when one assumes that every function in the representor satisfies the constraint $P(H_1 \wedge E) > P(H_2 \wedge E)$, it is still the case that updating on E will leave the agent’s imprecise credence in H_1 unchanged at the original value of $(0, 1)$.¹⁶ So even when the imprecise model helps itself to the crucial assumption at play in the comparative conditionalisation model, it still entails the impossibility of even the simplest inductive learning. This being the case, the advocate of the imprecise model clearly cannot claim that the progress of the comparative model presented here is paid for entirely by extra assumptions about the structure of the agent’s prior epistemic state. Even if those assumptions are integrated into the imprecise model, inductive learning remains impossible.

6 Conclusion

To summarise, the central results of the preceding analysis are that

- (i) There are two natural generalisations of the principle of indifference to the comparative setting, and although both yield consistent advice that can be consistently applied across multiple partitions, only CPI_2 provides *good* advice to agents at the beginning of their epistemic lives.
- (ii) In light of (i), it makes sense for those epistemologists that hope to offer advice to agents at the beginning of their epistemic lives to adopt CPI_2 as a principle of ideal rationality, which requires that they drop the Opinionation assumption as a constraint on the epistemic states

¹⁵This amended version of the imprecise model is also considered by i.e. Rinard (2013).

¹⁶To see this, note that Rinard’s (2013) derivation of the imprecise posterior credence in H_1 does not assume anything about the relative probabilities of the conjunctions $E \wedge H_1$ and $E \wedge H_2$. It relies only on the facts that (i) the prior credence in H_1 is maximally imprecise (either $[0, 1]$ or $(0, 1)$), and (ii) the fact that $P(E|H_1) = 1$ and $P(E|H_2) = \frac{1}{10}$, for every function P in the representor.

of rational agents.

- (iii) CPI_2 (unlike the imprecise approach to indifference) can be straightforwardly integrated into a plausible inductive logic that directly generalises Bayesian conditionalisation to the non-opinionated setting.

As Norton writes,

In one ideal, a logic of induction would provide us with a belief state representing total ignorance that would evolve towards different belief states as new evidence is learned.

That the Bayesian system cannot be such a logic follows from well-known, elementary considerations. (Norton, 2008: 45)

Here, I have shown that an inductive logic articulated in terms of non-opinionated confidence orderings is able to identify a unique and intuitively rational state of total ignorance that can evolve (in a plausible way) towards different belief states as evidence is learned. Norton's ideal is attainable.

Acknowledgements

I am grateful to Jeremy Butterfield, Lee Elkin, Stephan Hartmann, John Norton, Reuben Stern, and audiences in Munich and Pittsburgh for helpful feedback on earlier versions of this material. I'm also indebted to two anonymous referees for detailed and constructive replies to earlier drafts of this paper.

This work was supported by the Alexander von Humboldt foundation.

References

Alon, S. and Lehrer, E. (2014). Subjective multi-prior probability: A representation of a partial likelihood relation. *Journal of Economic Theory*, 151: 476-492

- Carnap, R. (1936). Testability and Meaning (Part I). *Philosophy of Science*, 3(4): 419-471.
- de Finetti, B. (1937). La prevision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincare*, 7: 1-68.
- Dupin, Y.B. (2015). Blurring Out Cosmic Puzzles. *Philosophy of Science* 82:879-891.
- Fine, T. L. (1973). *Theories of Probability*. Academic Press.
- Fishburn, P. C. (1986). The axioms of subjective probability. *Statistical Science* 1(3), 335-345.
- Fitelson, B. and Mccarthy, D. (unpublished). *Coherence*. Manuscript.
- Forrest, P. (1989). The Problem of Representing Incompletely Ordered Doxastic Systems. *Synthese*, 79:279-303.
- Halpern, J. (2003). *Reasoning About Uncertainty*. MIT Press.
- Harrison-Trainor, M., Holliday, W. and Icard, P. (2016). A note on cancellation axioms for comparative probability. *Theory and Decision*, 80: 159-166.
- Icard, T. (2016). Pragmatic Considerations on Comparative Probability. *Philosophy of Science*, 83: 348-370.
- Joyce, J. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives* 19(1): 153-178
- Joyce, J. (2010). A Defence of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives* 24(1): 281-323
- Kass, R. and Wasserman, L. (1996). The Selection of Prior Distributions by Formal Rules. *Journal of the American Statistical Association* 91: 1343-1370
- Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Keynes, J.M. (1921). *A Treatise of Probability*. New York: AMS.
- Koopman, B.O. (1940). The axioms and algebra of intuitive probability. *Annals of Mathematics*, 41(2): 269-292.
- Norton, J. (2008). Ignorance and Indifference. *Philosophy of Science* 75: 45-68

- Norton, J. (2010). Cosmic Confusions: Not Supporting Versus Supporting Not. *Philosophy of Science* 77: 501-523
- Pettigrew, R. (2016). Accuracy, Risk and the Principle of Indifference. *Philosophy and Phenomenological Research*, 92(1): 35-59
- Rinard, S. (2013). Against Radical Credal Imprecision. *Thought: a Journal of Philosophy*, 2(1): 157-165
- Rinard, S. (2014). The Principle of Indifference and Imprecise Probability. *Thought: a Journal of Philosophy*, 3(2): 110-114
- Shackel, N. (2007). Bertrand's Paradox and the Principle of Indifference. *Philosophy of Science*, 74: 150-175.
- Van Fraassen, B. (1989). *Laws and Symmetry*. Oxford: Clarendon Press.
- Von Mises, R. (1951). *Probability, Truth and Statistics*. New York: Dover.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.
- White, R. (2009). Evidential Symmetry and Mushy Credence. *Oxford Studies in Epistemology*, 3: 161-186.
- Weatherson, B. (2007). The Bayesian and the Dogmatist. *Proceedings of the Aristotelian Society* 107: 169-185
- Williamson, J. (2007). Motivating Objective Bayesianism: From Empirical Constraints to Objective Probabilities. In Harper, W. and Wheeler, G., editors, *Probability and Inference: Essays in Honor of Henry E. Kyburg Jr.*, 155-183. London: College Publications.