

The Bayesian Who Knew Too Much*

Yann Benétreau-Dupin[†]

Preprint, December 2014. Forthcoming in *Synthese*.

Abstract

In several papers, John Norton has argued that Bayesianism cannot handle ignorance adequately due to its inability to distinguish between neutral and disconfirming evidence. He argued that this inability sows confusion in, e.g., anthropic reasoning in cosmology or the Doomsday argument, by allowing one to draw unwarranted conclusions from a *lack of knowledge*. Norton has suggested criteria for a candidate for representation of neutral support. Imprecise credences (families of credal probability functions) constitute a Bayesian-friendly framework that allows us to avoid inadequate neutral priors and better handle ignorance. The imprecise model generally agrees with Norton's representation of ignorance but requires that his criterion of self-duality be reformulated or abandoned.

Keywords: Imprecise credence; ignorance; indifference; principle of indifference; Doomsday argument; anthropic reasoning

1 Introduction: Bayesian Reasoning and Ignorance

If you ask me what the probability is of rolling a 2 with a throw of a fair, cubic die, I will answer, " $\frac{1}{6}$." Now imagine that we are having a conversation about Shakespeare's play *The Tempest* and you ask me to assign a probability to the claim that the play has

*The final publication will be available at Springer via <http://dx.doi.org/10.1007/s11229-014-0647-3>. The accepted version of this paper is available at http://publish.uwo.ca/~ybenetre/Research_files/BWKTm.Preprint.pdf. Section 4 is an excerpt of (Benétreau-Dupin, *forth.*).

[†]Dept of Philosophy & Rotman Institute of Philosophy, Western University, London, Ontario, Canada. ybenetre@uwo.ca

16,633 words. My response would be that I simply do not have a clue. My knowledge of Elizabethan drama would allow me to say that the value falls within some interval, in which all values seem equally plausible. However, I would not be able to assign a probability to the claim that the play has *exactly* 16,633 words. If you insist on a probability, I would respond—echoing David Albert—by asking, “What part of ‘I don’t have a clue’ do you not understand?”

The simplest version of Bayesianism is ill-equipped to handle such a case. Reasoning with the principle of indifference, as in the case of the die, motivates a uniform probability distribution over some interval of values; if there are 20,000 equiprobable numbers in the plausible range of number of words for that play, each number will be assigned a probability of $\frac{1}{20,000}$. Thus according to usual Bayesian confirmation theory, learning the exact length of *The Tempest* is equivalent to confirming a claim with a very low probability. In contrast, suppose that we are discussing whether or not the 1956 movie *Forbidden Planet* is a good adaptation of *The Tempest*, and you ask me to assign a probability to the claim that Shakespeare himself authored that movie’s screenplay. I would *not* be indifferent about that claim; I would assign to it a very low probability, and I would be very surprised if it were confirmed. Yet orthodox Bayesianism does not allow one to distinguish these two confirmations: it sees them both as the confirmation of a low-probability claim, thereby justifying in both cases a same sense of surprise. One may legitimately be surprised by the confirmation of a low-probability proposition, but not by the confirmation of something about which we are entirely *indifferent*. In the absence of a representation of *neutral* degree of belief, we have no choice but to treat large numbers of alternatives about which we are equally *indifferent* as if they all were *improbable* propositions. Consequently, any hypothesis—however unlikely—that lends more support to the observed value than was initially given by the constant probability distribution will be seen as significantly confirmed.

Steven [Weinberg \(1987\)](#) resorted to a similar kind of confirmation when he appealed to anthropic constraints (i.e., constraints related to the possibility for life to exist) in order to explain the value of the cosmological constant (i.e., the vacuum energy density ρ_V), left indeterminate by existing theories. Anthropic considerations provide bounds on ρ_V , which can be neither too large (because then galaxies could not form) nor too small (because then observers would not have time to arise before the universe recollapses). However, the anthropically allowed range is quite large, and, in the absence of further theories, it leaves us with many equipossible values for ρ_V . Weinberg further argued that

anthropic considerations may have a stronger, predictive role. The idea is that we should conditionalize the probability of different values of ρ_V on the number of observers they allow: the most likely value of ρ_V is the one that allows for the largest number of galaxies (taken as a proxy for the number of observers).¹ The probability measure for ρ_V is then as follows:

$$dp(\rho_V) = \nu(\rho_V) \cdot p_*(\rho_V) d\rho_V,$$

where $p_*(\rho) d\rho_V$ is the prior probability distribution, and $\nu(\rho_V)$ the average number of galaxies which form for ρ_V . By assuming that there is no known reason why the likelihood of ρ_V should be special at the observed value, and because the allowed range of ρ_V is very far from what we would expect from available theories, Weinberg argued that it is reasonable to assume that the prior probability distribution is constant within the anthropically allowed range, so that $dp(\rho_V)$ can be calculated as proportional to $\nu(\rho_V) d\rho_V$ (Weinberg, 2000, 2). Moreover, by assuming that we are typical observers (and thereby adopting what Alexander Vilenkin (1995) called a ‘principle of mediocrity’), Weinberg predicted that the observed value of this parameter should be close to the mean of the anthropically allowed values. The initial uniform distribution is turned into a prediction, a sharply peaked distribution around a preferred value. Later observations of the cosmological constant have been taken to vindicate this reasoning (see, e.g., Weinberg, 2007).²

John Norton (2010) criticized the probabilistic representation of ignorance on which such arguments rest. He objects to the claims that 1) there can be a probability distribution over admissible values of some physical parameters when in fact those are left *indeterminate* by existing theories, and more strenuously that 2) the observed value has *low probability* instead of a *neutral* probability. Norton has suggested that the inability to distinguish between neutral and disconfirming evidence is the sign a fundamental flaw of Bayesianism, originating from the requirement of additivity. Therefore, for him, only a radically different inductive logic can adequately represent ignorance.

Norton’s challenge is valid against the Bayesian who knows too much (i.e., who represents ignorance or indifference with a single uniform probability distribution). But there are other Bayesians.

¹This assumption is contentious (see, e.g., (Aguirre, 2001) for an alternative proposal).

²The median value of the distribution obtained by such anthropic prediction is about 20 times the observed value ρ_V^{obs} , whereas predictions based on existing theories are 120 orders of magnitude higher than the observed value (Pogosian et al., 2004).

2 A Bayesian Failure?

2.1 Neutral vs. Disconfirming

Before I discuss in more detail what a representation of neutral degree of belief—i.e., a representation of indifference or ignorance—could look like, let me rephrase the problem under consideration in Bayesian terms.

The Bayesian approach to epistemology can be characterized as follows (see, e.g., [Joyce, 2010](#), 281-282):

- Belief is not all or nothing. One can assign degrees of belief to propositions.
- These degrees obey the laws of probabilities.³
- Learning implies updating an initial degree of belief (called prior) to obtain a posterior. Updating a prior for hypothesis H after acquiring evidence E will involve taking conditional probabilities and applying Bayes's theorem:

$$p(H|E) = \frac{p(H) \cdot p(E|H)}{p(E)},$$

where $p(E|H)$ denotes the probability of E conditional on H (i.e., the probability of E given H).

- Rational agents use their graded beliefs to choose actions with higher expected value.

A controversial but usual element of probability theory, the principle of indifference claims that one must assign the same probability value to equipossible events, or events about which we are equally ignorant or indifferent. If there are many such events, additivity dictates that each of them be assigned a uniform, low probability value, which is equivalent to saying that each event is improbable (or, equivalently, that their negation is probable).

We can now see how to rephrase the problematic arguments considered in §1:

1. The value of a parameter k is left indeterminate by our background knowledge. According to the principle of indifference, that indeterminateness is represented by a constant probability distribution widely spread over the admissible values of k , each of which having low probability $p(k|B)$.

³(1) For a probability function p , $\forall \alpha, p(\alpha) \geq 0$; (2) if α is logically true, then $p(\alpha) = 1$; (3) additivity: if α, β are incompatible ($p(\alpha \& \beta) = 0$), then $p(\alpha \vee \beta) = p(\alpha) + p(\beta)$. It follows from these laws that $\forall \alpha, p(\alpha) + p(\neg \alpha) = 1$.

2. A theory T makes the observed value k_{obs} much more probable:

$$p(k_{\text{obs}}|T\&B) \gg p(k_{\text{obs}}|B).$$

3. According to Bayes's theorem, we then have

$$\frac{p(T|k_{\text{obs}}\&B)}{p(T|B)} = \frac{p(k_{\text{obs}}|T\&B)}{p(k_{\text{obs}}|B)} \gg 1.$$

In other words, observing k_{obs} lends strong support to T .

For Norton, this confirmation is unwarranted. It is based on a flaw of Bayesianism itself—namely, it has no ability to represent neutrality due to additivity (Norton, 2010, 501-502). Assigning a definite low probability value to a proposition about which we are ignorant is then *turning ignorance into improbability*. It comes down to conflating disconfirming evidence ($p(H|E) \ll 1$) with neutral evidential support ($p(H|E) = p(\neg H|E)$).

2.2 A Non-Bayesian Notion of Neutral Support?

Norton (2007a, 2008, 2010) introduced the following criteria for a candidate for representation of neutral evidential support (or indifference, or ignorance):

- it cannot be additive (and therefore does not obey the laws of probability),
- it cannot be represented by the degrees of a one-dimensional continuum, such as the reals in $[0, 1]$,
- it must be invariant under redescription,⁴
- it must be invariant under negation: if we are ignorant or indifferent as to whether or not α , we must be equally ignorant as to whether or not $\neg\alpha$.

It is clear that usual Bayesianism cannot meet all these criteria. We should then look for another framework for an inductive logic that would allow one not only to express ignorance and indifference, but also to compare credences and carry out inferences and confirmation.

⁴The invariance under redescription only requires that the probability value that corresponds to neutral support for a same event must not depend on how this event is described. For instance, in the example given above in §1, book length was given in terms of number of words and could be redescriptioned in terms of number of pages or lines.

Norton confesses that he “know[s] of no adequate theoretical representation” of such a framework (Norton, 2010, 504). He then simply refers to this representation of neutral support as ‘ I ’, for ‘indifference’ or ‘ignorance’, with the following properties expressed in terms of a (non-probabilistic) credal function p :

- $\forall \alpha, p(\alpha) = I \rightarrow p(-\alpha) = I$ (invariance under negation),
- $\forall \alpha_1, \alpha_2$ mutually exclusive (but $\alpha_1 \vee \alpha_2 \not\equiv \top$), $p(\alpha_1) = p(\alpha_2) = I \rightarrow p(\alpha_1 \vee \alpha_2) = I$ (non-additivity).⁵

This framework must also preserve the values $p(\top) = 1 = 1 - p(\perp)$.⁶

This set of criteria for a representation of neutral evidential support is compelling. Because of this conflict with additivity, representing neutrality is a serious challenge for Bayesianism. But I will show that we need not abandon Bayesianism altogether, and that enriched versions of it already in use can satisfy these criteria to a certain extent.

3 A Bayesian Notion of Neutral Support

3.1 Bayesian Credences Need Not Have Sharp Values

It has been argued (see, e.g., Levi, 1974; Walley, 1991; Joyce, 2010) that Bayesian credences need not have sharp values, and that there can be imprecise probabilities.⁷ The difficulty to assign sharp values to credences was already raised by Kyburg (1978), who saw this as psychologically unrealistic. An imprecise probabilities model recognizes “that our beliefs should not be any more definitive or unambiguous than the evidence we have for them.” (Joyce, 2010, 320)

It is possible to reject a precise Bayesian model in favor of a less exact one, in which credences are not well-defined but allow for *imprecise values*. Joyce defended an imprecise model in which credences are not represented merely by a range of values, but rather by a *family* of (probabilistic) credence functions. In this imprecise probability model,

⁵Strictly speaking, it is not entirely appropriate to define this condition in terms of additivity. For a representation of credence to be ‘non-additive’ in the sense of interest to Norton here, it has to fulfill the following condition: $\forall \alpha, \beta$ incompatible propositions about which we are completely indifferent or ignorant, we can have neither $p(\alpha \vee \beta) > p(\alpha)$ nor $p(\alpha \vee \beta) > p(\beta)$.

⁶ \top is an unconditionally true statement, and \perp an unconditionally false one.

⁷‘Imprecise credence’ is more appropriate than ‘imprecise probability’ since it does not necessarily obey the laws of probability. Here I nevertheless use both expressions interchangeably, as is done in the literature.

1. a believer's overall credal state can be represented by a family C of credence functions $[c_i]$ (...). Facts about the person's opinions correspond to properties common to all the credence functions in her credal state.
2. If the believer is rational, then every credence function in C is a probability.
3. If a person in credal state C learns that some event D obtains (...), then her post-learning state will be $C_D = \left\{ c(X|D) = c(X) \frac{c(D|X)}{c(D)} : c \in C \right\}$.
4. A rational decision-maker with credal state C is obliged to prefer one action α to another α^* when α 's expected utility exceeds that of α^* relative to every credence function in C . (Joyce, 2010, 288)

In other words, in this imprecise probability model, a credal state can be represented by a family of functions that behave as usual Bayesian, probabilistic credence functions. Joyce also offers an analogy that illustrate this model: the overall credal state C acts as a committee whose members (each being analogous to a credence function c_i) are rational agents who do not all agree with each other and who all update their credence in the same way, by conditionalizing on evidence they all agree upon. With this analogy, the properties of the overall credal state C correspond to those common to all the committee members.

There are several criteria for decision-making with imprecise probabilities between two propositions. Depending on the criterion chosen, one will prefer an event to another event if

- it has maximum lower expected value (Γ -minimax criterion),
 - it has maximum higher expected value (Γ -maximax),
 - it has maximum expected value for all distributions in the credal set (maximality),
 - it has a higher expected value for at least one distribution in the credal set (E -admissibility),
- or
- its lower expected value on all distributions in the credal set is greater than the other event's highest expected value on all distributions (interval dominance).⁸

⁸This list is not exhaustive. It is beyond the scope of this paper to compare and assess these criteria. See (Troffaes, 2007; Augustin et al., 2014, §8) for reviews.

This model allows one to simultaneously represent sharp and imprecise credences, but also comparative probabilities. It can accommodate sharp credences, for which the usual condition of additivity holds. But it can also accommodate less sharply defined relationships when credences are fuzzy. It does so by means of a family of credence functions, each of which is a Bayesian function that obeys the laws of probability.

3.2 Neutral Support with Imprecise Credences

This imprecise model allows one to distinguish two notions of neutral support: distinguish stochastic independence (true independence of evidence for a given hypothesis) and unknown interaction (unknown dependence).

With a single (precise) probability distribution, two variables α_m and α_n in appropriate algebras are stochastically independent if $p(\alpha_m|\alpha_n) = p(\alpha_m)$ whenever $p(\alpha_n) > 0$. Different concepts have been suggested to extend the notion of stochastic independence of two variables α_m and α_n to the imprecise model (see [Cozman, 2012](#)), among which:

- complete independence, if stochastic independence of α_m and α_n obtains for each distribution $c_i \in C$,⁹
- confirmational irrelevance, if $C(\alpha_m|\alpha_n) = C(\alpha_m)$,
- epistemic irrelevance, if stochastic independence obtains for the lowest expectation among all functions $c_i \in C$, or the related, symmetric concept of epistemic independence (if α_m is epistemically irrelevant to α_n and α_n to α_m).

In contrast, unknown interaction between two variables α_m and α_n can be represented by a credal set that contains credence functions that differ on the correlation between α_m and α_n . Such a credal set would include credence functions c_i that are such that $c_i(\alpha_m \& \alpha_n) \leq c_i(\alpha_m) \cdot c_i(\alpha_n)$ and others functions c_j such that $c_j(\alpha_m \& \alpha_n) > c_j(\alpha_m) \cdot c_j(\alpha_n)$. Following the committee analogy once again, in a state of unknown interaction between α_m and α_n , jury members disagree with each other about that interaction. The overall credal state of the committee neither favors nor dismisses any of the opinions of its jury members; it can only express a lack of agreement.

The case of complete ignorance or indifference is of particular interest for the anthropic argument presented in §1: we want to represent *complete ignorance or indifference* as to

⁹As discussed in ([Cozman, 2012](#)), this definition violates convexity.

what value a certain physical parameter should have. With the imprecise model, recalling the committee analogy, complete ignorance as to whether or not α (i.e., complete indifference between which of α or $\neg\alpha$ is more likely) can be represented by a committee in which there is *no agreement* among its members about whether or not α is more or less probable than $\neg\alpha$. This notion of neutral support is analog to Norton’s requirement that, in order to express neutral support for an event α , a measure of belief about α be equal to a corresponding measure of *disbelief* about α (Norton, 2007a). In the context of credal sets, this requirement is fulfilled, for example, by a committee composed as follows:

for every jury member and for every contingent event α , there exists a jury member who has as much disbelief in α as another has belief in it.

Or, to put it in terms of credence functions,

if C is a credal state representing a family of credence functions c about contingent events α , if $\forall c, \alpha, \exists c' \in C$ such that $c'(\alpha) = 1 - c(\alpha)$, then C corresponds to a state of ignorance.¹⁰

If all the jury members are rational (i.e., if all the credence functions are probability functions), we must have, by definition, $C(\top) = 1 = 1 - C(\perp)$.¹¹

Such a credal state about a given event α can then be updated by conditionalizing upon new evidence, taken as such by all the jury members (i.e., all the credence functions $c \in C$). Thereby, after new evidence E is gathered, the range of values taken by all the credence functions of C , $c(\alpha|E)$, is susceptible to change, and so is the overall credal state regarding α .

A trivial but extreme example of representation of *complete* ignorance by means of a set of credal functions is the set \mathfrak{J} of *all* possible probability distributions. For any proposition α about which are completely ignorant, that representation of ignorance would give us $C(\alpha) = [0, 1]$. In case of complete ignorance, excluding possible probability distributions compatible with our evidence is “pretending to have information [we do] not possess.” (Joyce, 2005, 170) This proposal meets all of Norton’s criteria for a representation of ignorance (see *infra*, § 2.2).

¹⁰It may be surprising to read that, for a certain proposition, a credal set that gives the set of values $\{0.49, 0.51\}$ is a better representation of ignorance than one that gives the set of values $\{0.1, 0.8\}$. It is here a representation of ignorance that gives as much support to a proposition as it does to its negation that this criterion captures.

¹¹i.e., $\forall c \in C, c(\top) = 1 = 1 - c(\perp)$.

There is however a good reason not to be content with such an extreme representation of ignorance. Indeed, in that set \mathfrak{J} of all possible probability distributions will be extremely sharp probability distributions that require an unreasonably large—or even infinite—number of updatings before they can yield posteriors distributions that are significantly different (Rinard, 2013). Such distributions in \mathfrak{J} are said to be dogmatic, and consequently the whole set \mathfrak{J} is dogmatic. A representation of complete ignorance \mathfrak{J} , and generally any vacuous prior, entails a vacuous posterior. This should prevent such a set from being used in an inferential process in which we may hope to move *away* from a state of ignorance after a certain number of iterations of Bayesian updating. This representation of ignorance by means of a family of credal functions, although it satisfies Norton’s criteria for ignorance, is *incompatible with learning*. That is why imprecise statisticians, who are interested in inferential processes, prefer to deal with ‘near-ignorance’ (i.e., broad credal intervals smaller than $[0,1]$) rather than complete ignorance, thereby ruling out dogmatic priors (see Moral, 2012; Augustin et al., 2014; Walley, 1991, §7.3.7).

By way of example, let us represent our near-ignorance about three mutually exclusive propositions $\alpha_1, \alpha_2, \alpha_3$ with a credal set $C = \{c_i\}$ consisting of probability functions defined as follows:

	α_1	α_2	α_3
$c_1(\alpha_i)$	0.9	0.05	0.05
$c_2(\alpha_i)$	0.05	0.9	0.05
$c_3(\alpha_i)$	0.05	0.05	0.9

and for all $i, j \in \{1, 2, 3\}$, for all $x \neq \alpha_i$ or their disjunctions, $c_j(x) = 0$.

In this credal set, no value of α_m is preferred to the others (i.e., no value of α_m is favored by all three distributions $c_i \in C$); and for all $\alpha_m, \alpha_n, m \neq n, C(\alpha_m) = C(\alpha_n)$. A constant, unique probability distribution could express this as well. But this credal set tells us more than that, namely that, for all m , no proposition α_m is preferred to its negation (i.e., no value of α_m is favored over its negation by all three distributions $c_i \in C$).

This example does not satisfy the characterization of ignorance proposed earlier. However, we can see that, with a few amendments, this set suffices to represent indifference about any of these propositions α_i , and can meet the requirements for a representation of neutral support (see above §2.2):

- it is not a sharp value in $[0, 1]$,

- it can be defined so as to be invariant under redescription: $\forall \alpha, a$ (a , redescription of α), if $C(\alpha)$ represents a state of ignorance, so will $C(a)$,¹²
- it is invariant under negation (no proposition α is preferred to its negation).
- we do have $C(\alpha_1 \vee \alpha_2 \vee \alpha_3) = C(\top) = 1 = 1 - C(\perp)$.

The criterion of non-additivity (see *infra*, note 5) cannot be satisfied in a trivial manner. In the example above, we have $\forall i, m, n \in \{1, 2, 3\}, m \neq n, c_i(\alpha_m \vee \alpha_n) > c_i(\alpha_m)$. That is also true of the bounds of the credal intervals (they will not be non-additive). But we can prevent that by adding to our set credal functions c_k such that $c_k(\alpha_1 \vee \alpha_2) = c_k(\alpha_1)$ and so on. That is, we can add to our initial set three other functions c_4, c_5 , and c_6 so as to have $\forall i \in \{1, 2, 3\}, \exists c_j \in C, j \in \{4, 5, 6\}$ such that $c_j(\alpha_i) = 0$. But the newly added functions are not reasonable, since if we agreed that the propositions α_i are contingent, then no function ruling them out completely should be accepted in our credal set.

However, adopting certain conventions can mitigate the effects of additivity. If among the criteria for decision-making with credal sets introduced earlier in § 3.1 we choose that of interval dominance, then our representation of credence is non-additive in this example (i.e., for all $m, n \in \{1, 2, 3\}, \alpha_m \vee \alpha_n$ is not preferred to α_m). However, one might argue that interval dominance is often not a desirable criterion. It is a demanding criterion that allows one to express a preference between two propositions only if one is unambiguously better than the other. This criterion is arguably not fine-grained enough to help us for most of the inferences we are likely to encounter. Other, often more desirable decision-making criteria, however, will result in additive imprecise credences. With other decision rules, we could under circumstances circumvent additivity by adopting threshold values, beyond (respectively below) which all values are considered to be equally confirmatory (respectively disconfirmatory). For instance, in the set given above, if we consider that any value below 0.1 is equally disconfirmatory and any value beyond 0.9 is equally confirmatory, then we lose all additivity.

We could also adopt a different strategy that does not involve such conventions. For any proposition α about which we are ignorant (whether the α_i or their Boolean combinations), we can define a credal set C such that $C(\alpha) = C(\neg\alpha)$. The imprecise model allows us to treat equipossible propositions in the same way, which does not mean that they must

¹²If the functions in this set are described as Dirichlet distributions, then this criterion will be satisfied (see, e.g., [de Cooman et al., 2009](#)).

receive a same probability value or be represented by the same credal set. By ‘sameness of treatment’ of contingent, equipossible events about which we are ignorant or indifferent, I mean that our credal state of ignorance about them would be modeled in the same manner, by a credal set having the same desired properties. We will come back to this in the next section.

With the imprecise model, not one single set of functions or one set of rules to define it will be suited to truly represent ignorance in all situations, unless we are ready to represent a state of complete ignorance by the undesirable and unreasonable set \mathfrak{F} of all possible credal functions. But this model allows one to define sets that do not favor any of the propositions about which we are ignorant and that is suited to a particular ensemble of propositions under consideration. It does so by means of Bayesian credence functions, which allows for our credal state to be updated and evolve. All that this requires is that we do not demand that agent’s credences have sharp values.

3.3 Norton’s Objections

3.3.1 Interpretative objections

Norton (2007a,b, 2008) has formulated several objections to the imprecise model. He has expressed a general discomfort with what he considers to be an inadequate approach, “an attempt to simulate an inherently nonadditive logic with an additive measure, rather than to seek the logic directly.” (Norton, 2010, 504, note 4). But as indirect or contrived as this method might seem, it is successfully applied in statistical analysis (see, e.g., Walley, 1991; Augustin et al., 2014). Yet Norton has raised more pointed criticisms of the imprecise model, about the very question of representing ignorance:

the representation [of ignorance by sets of probability functions] is not literally correct. That is, ignorance is not the maintaining of all possible beliefs at once; it is the maintaining of none of them. So we should regard the device of convex sets as a way of simulating ignorance through a convenient fiction.¹³ (Norton, 2007b, §4.2)

Or elsewhere:

¹³This remark also applies to non-convex sets.

The sort of ignorance I seek to characterize is first order ignorance; it is just not knowing which is the true outcome; not a second order uncertainty about an uncertainty. (Norton, 2007a, §6.2)

Such criticism is in fact not specific to the question of representing ignorance. It is aimed more generally at the use of *several* credal functions in order to represent a *unique* credal state.¹⁴ But it does not apply to the imprecise model we have considered here, and, in general, proponents of imprecise probabilities need not endorse this view. Indeed, this model does not allow “the maintaining of all possible beliefs at once.” Even though a credal set may be comprised of several credal functions, the agent’s credence it represents is unique; its properties are those that are common to all the credal functions in that set. If all possible beliefs could be held at once—or rather, if no particular belief can be preferred to any other—an agent’s credence would not be multiple, it would just *not be any* of these particular beliefs.

3.3.2 Indifference and self-duality

Another criticism more specifically aimed at the issue of representing ignorance or indifference deserves a closer examination. Norton (2007a) argues that a representation of ignorance should satisfy what he calls a condition of self-duality. The dual of a measure of belief is a measure of disbelief. If we are ignorant about a proposition α , our degree of belief that α should not be different from our degree of disbelief that α : “the state [of complete ignorance] is self-dual in its contingent propositions (...).” (Norton, 2007a, §6.1, 247) This follows from the requirement of invariance under negation for a representation of neutral support.

We saw earlier in § 3.2 that no single probability function (or in general no measure, additive by definition) can meet this criterion.¹⁵ Any probability function will necessarily express a (non-strictly) increasing belief as we go from propositions to their logical consequences (e.g., for any functions c in a credal set, $\alpha_1 \vdash \alpha_1 \vee \alpha_2 \rightarrow c(\alpha_1 \vee \alpha_2) \geq c(\alpha_1)$). On the other hand, a measure of disbelief should be (non-strictly) *decreasing* as we go

¹⁴The following passage makes it clear that Norton thinks of the use of a set of probability functions as allowing the simultaneous representation of several states of belief: “the use of sets renders ignorance as a second order sort of belief. We allow that many different belief-disbelief states are possible. We represent ignorance by presenting them all, in effect saying that we don’t know which is the pertinent one.” (Norton, 2007a, §6.2, 248)

¹⁵Unless we are dealing with only two mutually exclusive propositions.

from propositions to their logical consequences. Consequently, no probability function can simultaneously express our belief and our disbelief in a certain proposition, even if we consider only contingent propositions. This, however, is not necessarily true of credal sets. We also saw in § 3.2 that if decision-making is based on interval dominance, then imprecise credences need not be additive for contingent propositions. Consequently, an imprecise credence represented by a set of functions as the one given in § 3.2 can simultaneously represent belief and disbelief (at least for contingent propositions).

On Norton's account (Norton, 2007a, § 3.2,6.2), a credal set C representing ignorance is self-dual if it contains probability functions c_i and their duals c'_i such that $\forall c_i \in C, c'_i$ is such that $\forall \alpha, c_i(\alpha) = c'_i(\neg\alpha)$ and $c_i(\neg\alpha) = c'_i(\alpha)$. It is clear that we cannot have such a set. Indeed, if we have

$$\forall c \in C, \exists c' \in C \text{ such that } \forall \alpha, c'(\alpha) = 1 - c(\alpha),$$

then if $\alpha \equiv \top$ and c is a probability function, we would get

$$c'(\alpha \equiv \top) = 1 - c(\top) = 0,$$

and then c' would break the second law of probability (see *infra* note 3). Norton restricts the requirement of self-duality to contingent events only, which are the only relevant ones when we want to define a representation of ignorance or indifference. But even if we forget about logical truths and contradictions, additivity will forbid such a self-dual set. Indeed, for two propositions α and β mutually exclusive, we do not get $\forall i, c'_i(\alpha \vee \beta) = 1 - c_i(\alpha \vee \beta) = c'_i(\neg(\alpha \vee \beta))$. If it were the case, we would obtain

$$\begin{aligned} c'_i(\alpha \vee \beta) &= c_i(\neg(\alpha \vee \beta)) \\ &= c'_i(\alpha) + c'_i(\beta) \\ &= c_i(\neg\alpha) + c_i(\neg\beta) \end{aligned}$$

For i such that $c_i(\alpha) \ll 1$, $c_i(\neg\alpha) \simeq 1$ and we may have $c'_i(\alpha \vee \beta) = c_i(\neg\alpha) + c_i(\neg\beta) > 1$, which would prevent c'_i from being a probability function. If this set is not self-dual for α and β mutually exclusive, it cannot be self-dual for all α and the credal state as I have defined it cannot be self-dual.

In contrast, as we saw above in § 3.2, for any set of contingent propositions, the imprecise model allows us to define sets that represent our indifference or ignorance. The proposal for a representation of neutral support we saw above in § 3.2, according to which if C is

a credal state representing a family of credence functions c about contingent events α , C corresponds to a state of indifference if $\forall c, \alpha, \exists c' \in C$ such that $c'(\alpha) = 1 - c(\alpha)$, differs from Norton's criterion of self-duality (note the different placement of the universal quantifier on α). The imprecise model as I have considered it here does not guarantee us that once a credal set is defined so as to represent indifference about a set of propositions, it will necessarily adequately represent our indifference about all their Boolean combinations, nor, in general, that it will adequately provide neutral support to any other proposition we are indifferent about. This model does not suggest either that the same rules used to define a credal set representing indifference should apply to all situations (this will be illustrated below in § 4).

If we consider only contingent propositions, and if decision-making is based on interval dominance, then with imprecise probabilities we have at our disposal a representation of ignorance or indifference that shares with a self-dual representation the relevant properties for a representation of neutral support, namely the ability to simultaneously represent belief and disbelief. With other decision rules (e.g., Γ -minimax), a self-dual representation of ignorance or indifference could only express nothing less than *complete* ignorance. We saw in § 3.2 that such a credal set, \mathfrak{I} , exists, but we also saw that it is incompatible with learning. The less demanding representations of ignorance or indifference I have considered here have over a self-dual representation a clear expressive advantage. We saw that the imprecise model allows one to distinguish between stochastic independence, epistemic irrelevance, unknown interaction, or ignorance or indifference about the value of a parameter. It can fulfill criteria of a representation of ignorance better than what a single probability function can do, yet it does so by means of probability functions, each of which can subsequently be modified following Bayes's rule. Consider for instance the 3-function example from § 3.2. Now assume that the probability of α_i in that example corresponds to the probability of drawing the numbered ball B_i from an urn (with replacement between each draw). Assume further that we are interested in determining a bias—of our urn or the balls—in this drawing. How should we represent our initial state of ignorance in a way that does not prevent us from eventually finding a value after a sufficient iteration of Bayesian updating? It is unclear how Norton's non-probabilistic, self-dual representation of complete ignorance, ' I ', can evolve, whereas it is possible for our 3-function credal set to yield such a result eventually.¹⁶ Notwithstanding the lack of self-duality, the ability for

¹⁶See (Piatti et al., 2009; Moral, 2012) for a discussion about the conditions on a near-ignorance credal set required for learning. These requirements favor the use of Dirichlet distributions.

credal sets representing indifference or ignorance to be updated is a desirable feature that makes an imprecise representation of indifference a more interesting element of inductive logic than Norton’s self-dual measure I .

The kind of representation of ignorance that Norton seeks is part of a larger inductive logic yet to be carried out. It is in the light of the search for a unique representation of our credences that Norton imposes this criterion of self-duality of complete ignorance, applicable to *any* proposition about which we are ignorant. In this context, it is a plausible requirement. But it is a very strong condition that can leave one wondering how a unique representation as the one Norton proposes—which would yield the same value ‘ I ’ whatever the event about which we are ignorant—can be used in a fruitful inductive process. Norton offers no compelling reason why his demanding notion of self-duality should be required of alternative representations of ignorance or indifference such as the ones the imprecise model affords us.

4 Dissolving Cosmic Confusions

In order to emphasize how ill-equipped Bayesianism (or any other inductive framework based on additive measures) is in order to deal with ignorance, and to show its inability to avoid unwarranted conclusions, Norton (2010) invited us to consider two instances of ‘cosmic confusion.’ These are anthropic reasoning (see *infra*, § 1) and the Doomsday argument. In both cases, we draw conclusions from a *lack of knowledge*. We can see how the use of imprecise credences can dissolve what Norton referred to as ‘cosmic confusions’ in a way no unique probability distribution can.¹⁷

The Doomsday argument is a family of arguments about humanity’s likely survival (see, e.g., (Bostrom, 2002, §6-7), (Richmond, 2006) for reviews). It allows one to compare the likelihood of two scenarios about humanity’s survival or even make a prediction about the end date for humanity based only on the assumption that our place on humanity’s timeline is random. Some versions of this argument have been addressed within the framework of orthodox Bayesianism so as to block any conclusion we could draw from these assumptions alone. But a variant of this argument, based on the assumption that we have a random birth rank among all humans (Gott, 1994), cannot be dissolved without appealing to imprecise credences. The argument goes as follows. Let r be my birth rank (i.e., I am the r^{th} human

¹⁷I have given a more detailed presentation of the problems and solutions discussed in this section in (Ben treau-Dupin, *forth.*).

to be born), and N the total number of humans that will ever be born.

1. Assume that there is nothing special about my rank r . Following the principle of indifference, whatever r , the probability of r conditional on N is $p(r|N) = \frac{1}{N}$.
2. Assume the following improper prior¹⁸ for N : $p(N) = \frac{k}{N}$. k is a normalizing constant whose value does not matter.
3. This choice of distributions $p(r|N)$ and $p(N)$ gives us the prior distribution $p(r)$:

$$p(r) = \int_{N=r}^{N=\infty} p(r|N)p(N) dN = \int_{N=r}^{N=\infty} \frac{k}{N^2} dN = \frac{k}{r}.$$

4. Then, Bayes's theorem gives us

$$p(N|r) = \frac{p(r|N) \cdot p(N)}{p(r)} = \frac{r}{N^2},$$

which favors small N and allows us to make an estimate for N at any confidence-level.

This result should strike us as surprising: we should not be able to learn something from nothing! Indeed, according to that argument, we can make a prediction for N based only on knowing our rank r and on *not* knowing anything about the probability of r conditional on N , i.e., on being indifferent—or equally uncommitted—about any value it may take.

In this argument, any choice of prior probability distribution will result in a prediction for N , at any confidence-level. However, if our prior ignorance or indifference about N , $C(N)$, is represented by a credal set containing an *infinity* of credal functions, $\{c : c \in C\}$, each normalizable, defined on $\mathbb{N}^{>0}$, and such that $\forall c \in C, \lim_{N \rightarrow \infty} (c(N)) = 0$ (e.g., a family of Pareto distributions), then the resulting prediction for N diverges. In other words, this imprecise representation of prior credence in N , reflecting our ignorance about N , does not yield any prediction about N . Without the possibility for my prior credence to be represented not by a single probability distribution, but instead by an *infinite set* of probability distributions, I cannot avoid obtaining an arbitrarily precise prediction.

¹⁸As [Gott \(1994\)](#) recalls, this choice of prior is fairly standard (albeit contentious) in statistical analysis. It is the Jeffreys prior for the unbounded parameter N , such that $p(N) dN \propto d \ln N \propto \frac{dN}{N}$. This means that the probability for N to be in any logarithmic interval is the same. This prior is called improper because it is not normalizable, and it is usually argued that it is justified when it yields a normalizable posterior.

In the case of the cosmological constant problem (see *infra* § 1), representing our prior ignorance or indifference about the value of the vacuum energy density ρ_V by an imprecise credal set can limit, if not entirely dissolve, the appeal of anthropic considerations. As we saw earlier, [Weinberg \(1987\)](#) argued that, in the absence of useful theoretical background, it was reasonable to assume a constant, uniform prior probability distribution for ρ_V within the anthropically allowed range, and then conditionalize on the number of observers each value of ρ_V would allow for. With the imprecise model, a state of indifference between different values of ρ_V within the anthropically allowed range can be expressed by a set of probability distributions $\{c_\star : c_\star \in C_\star\}$, all of which normalizable over the anthropic range and such that $\forall \rho_V, \exists c_{\star i}, c_{\star j} \in C_\star$ such that ρ_V is favored by $c_{\star i}$ and not by $c_{\star j}$.¹⁹ It is in principle possible to define this prior credal set so that for any value of ρ_V , the lowest expectation value among the the posteriors is lower than the highest expectation value among the priors. If then we adopt interval dominance as a criterion for decision-making, then no observation of ρ_V will be able to lend support to our anthropic prediction.

As we saw earlier, one may object to the adoption of interval dominance in such a case. This choice of demanding decision rule could be motivated by the fact that we have no plausible alternative theoretical framework to the anthropic argument. In this context, it can be reasonable to agree to increase one's credence about the anthropic explanation *only if* it does better than any other yet unknown alternative might have done. Nonetheless, if we adopt other decision rules, it is possible with the imprecise model to construct prior credal sets that define a large interval over the anthropic range such that the confirmatory boost obtained after observing ρ_V is not nearly as vindicative as it is with a single, uniform distribution.

5 Conclusion

[Norton \(2010\)](#) has correctly argued that representing neutrality with a broadly spread single probability distribution amounts to conflating ignorance with improbability. He has shown how this leads to unwarranted confirmations. I here claim that a credal state of ignorance should best be represented by an imprecise credence. With this approach, merely acquiring information about the value of a certain parameter cannot suffice to justify a sense of surprise. The imprecise model offers us a more adequate representation of neutrality

¹⁹For reasons expressed earlier in footnotes 12 and 16, this should preferably be done by means of Dirichlet distributions.

and prevents prior credences from doing too much inductive work, as is illustrated by its ability to block the consequence of the Doomsday argument better than what orthodox Bayesianism can do.

We saw that, if we adopt interval dominance as a criterion for decision-making, it is possible for an imprecise representation of indifference to meet the criteria for a representation of neutral support put forth in (Norton, 2007a, 2008, 2010). But if we adopt less demanding decision-making rules, we can still construct a representation of indifference by means of credal sets that meet these criteria to a large extent, and we can do so without compromising Bayesianism altogether. It only requires that we do not demand that credences be sharp nor that a unique representation be applicable to all cases of ignorance or indifference (i.e., that self-duality be abandoned).

One can see Norton's argument as emphasizing the perils of excessive structure imported from probability theory into inductive logic, and then arguing that we need to eliminate much of that structure. There are several ways to modify the mathematical structure to counter the assumption of additivity. Norton's stated motivations are not sufficient to force us to adopt the framework he advocates. The imprecise model has the advantage of allowing us to distinguish different kinds of ignorance and indifference. More importantly, it makes it possible to move out of a state of ignorance when we acquire interesting information.

Acknowledgements

I am grateful to Wayne Myrvold for initial discussions, Jim Joyce and John Norton for stimulating exchanges. I am indebted to Chris Smeenk for many comments and suggestions.

References

- Aguirre, A. (2001, September). Cold Big-Bang Cosmology as a Counterexample to Several Anthropic Arguments. *Physical Review D* 64(8), 1–12.
- Augustin, T., F. P. Coolen, G. de Cooman, and M. C. Troffaes (Eds.) (2014). *Introduction to Imprecise Probabilities*. Wiley & Sons.
- Benétreau-Dupin, Y. (Forthcoming). Blurring Out Cosmic Puzzles. *Philosophy of Science*.

- Bostrom, N. (2002). *Anthropic Bias: Observation Selection Effects in Science and Philosophy*. Routledge.
- Cozman, F. G. (2012). Sets of Probability Distributions, Independence, and Convexity. *Synthese* 186(2), 577–600.
- de Cooman, G., E. Miranda, and E. Quaeghebeur (2009, February). Representation Insensitivity in Immediate Prediction under Exchangeability. *International Journal of Approximate Reasoning* 50(2), 204–216.
- Gott, J. R. (1994). Future Prospects Discussed. *Nature* 368(March), 108.
- Joyce, J. M. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives* 19, 153–178.
- Joyce, J. M. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives* 24(1), 281–323.
- Kyburg, H. (1978). Subjective Probability: Criticisms, Reflections, and Problems. *Journal of Philosophical Logic* 7(1), 157–180.
- Levi, I. (1974). On Indeterminate Probabilities. *The Journal of Philosophy* 71(13).
- Moral, S. (2012, April). Imprecise Probabilities for Representing Ignorance about a Parameter. *International Journal of Approximate Reasoning* 53(3), 347–362.
- Norton, J. D. (2007a, October). Disbelief as the Dual of Belief. *International Studies in the Philosophy of Science* 21(3), 231–252.
- Norton, J. D. (2007b, May). Probability Disassembled. *The British Journal for the Philosophy of Science* 58(2), 141–171.
- Norton, J. D. (2008). Ignorance and Indifference. *Philosophy of Science* 75(1), 45–68.
- Norton, J. D. (2010). Cosmic Confusions: Not Supporting versus Supporting Not. *Philosophy of Science* 77(4), 501–523.
- Piatti, A., M. Zaffalon, F. Trojani, and M. Hutter (2009, April). Limits of Learning about a Categorical Latent Variable under Prior Near-Ignorance. *International Journal of Approximate Reasoning* 50(4), 597–611.

-
- Pogosian, L., A. Vilenkin, and M. Tegmark (2004, July). Anthropic Predictions for Vacuum Energy and Neutrino Masses. *Journal of Cosmology and Astroparticle Physics* 7(005), 1–17.
- Richmond, A. (2006). The Doomsday Argument. *Philosophical Books* 47(2), 129–142.
- Rinard, S. (2013). Against Radical Credal Imprecision. *Thought: A Journal of Philosophy* 2(2), 157–165.
- Troffaes, M. (2007). Decision Making Under Uncertainty Using Imprecise Probabilities *International Journal of Approximate Reasoning* 45(1), 17–29.
- Vilenkin, A. (1995). Predictions from Quantum Cosmology. *Physical Review Letters* 74(6), 4–7.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.
- Weinberg, S. (1987). Anthropic Bound on the Cosmological Constant. *Physical Review Letters* 59(22), 2607–2610.
- Weinberg, S. (2000). A Priori Probability Distribution of the Cosmological Constant. *arXiv preprint astro-ph/0002387*, 0–15.
- Weinberg, S. (2007). Living in the Multiverse. In B. Carr (Ed.), *Universe or Multiverse?*, Chapter 2, pp. 29–42. Cambridge University Press.