Suspension of judgment, non-additivity, and additivity of possibilities

Aldo Filomeno*

Abstract

In situations where we ignore everything but the space of possibilities, we ought to suspend judgment—that is, remain agnostic—about which of these possibilities is the case. This means that we cannot sum our degrees of belief in different possibilities, something that has been formalized as an axiom of non-additivity. Consistent with this way of representing our ignorance, I defend a doxastic norm that recommends that we should nevertheless follow a certain additivity of possibilities: even if we cannot sum degrees of belief in different possibilities, we should be more confident in larger groups of possibilities. It is thus shown that, in the type of situation considered (in so-called 'classical ignorance', i.e. "behind a thin veil of ignorance"), it is epistemically rational for advocates of suspending judgment to endorse this comparative confidence; while on the other hand it is shown that, even in classical ignorance, no stronger belief—such as a precise uniform probability distribution—is warranted.

Keywords

Ignorance; Decision under uncertainty; Objective Bayesianism; Material theory of induction; Suspension of judgment; Agnosticism; Non-additivity.

*Instituto de Filosofía, Universidad Católica de Valparaíso.

Forthcoming in Acta Analytica

Contents

1 The norm

Imagine that in the Indiana Jones trilogy (let's pretend the next were never made) there is a scene in which Indiana Jones is in front of 10 chests and there is treasure inside one of them. Indiana ignores everything else. According to a widely-held view in epistemology (e.g., White 2005, de Cooman and Miranda 2007, Norton 2008, Friedman 2013a,b, Myrvold

2014, and Benétreau-Dupin 2015), if he is asked where the treasure is he should answer, "I do not know"; his doxastic attitude should be to *suspend judgment*, that is, to remain *agnostic*. This means that he should *not* hold a 0.1 degree of belief that the treasure is inside each of the chests—the traditional proposal of the objective Bayesian, based on the principle of indifference. He is in a state of ignorance, call it 'I', which is not a numerical degree of belief. The agnostic view holds that having the same numerical degree of belief in each possibility is not only *too precise* a judgment, it is a *judgment*, and as such it is unwarranted, since being in an epistemic state of ignorance means *suspending* judgment. This is not the only discrepancy, as we will see throughout the paper.[1]

Given that he is suspending judgment about where the treasure is, he is as ignorant about the treasure being in chest 1 as about the treasure being in, say, $\{$chest 5 $\cup$ chest 6 $\cup$ chest 7$\}$. Ignorance, it is said, is *not additive*. That is, we cannot sum the degrees of belief in different possibilities: in fact, it just might well be more likely (in the sense of objective chance) that the treasure is in chest 1, and in any case we do not even have numerical degrees of belief to sum.

(The axiom of additivity states that, for any mutually disjoint sets of possibilities A and B, $P(A \cup B) = P(A) + P(B)$; where $P(X)$ is the degree of belief in X being the case. As further explained in section 2, in situations of ignorance additivity is unwarranted. In its stead, an axiom of non-additivity states, roughly, that for any sets of possibilities A and B, $P(A) = I, P(B) = I$, and $P(A \cup B) = I$, where I stands for the non-numerical doxastic state of agnosticism.)

In this paper, we will assume this 'agnostic' way of representing our doxastic state of ignorance; that is, we will assume that in situations of ignorance we have to suspend

[1]For recent philosophical discussions on other aspects of suspension of judgment, see e.g. Friedman (2013a,b, 2015); Tang (2015); Staffel (2019); Rosa (2020); McGrath (2020); Filomeno (2022). Its origins date back to the ancient scepticism of Pyrrho (Empiricus, I c. A.C.). A number of formal representational frameworks in inductive logic have been elaborated to properly represent suspension of judgment. The different approaches include non-probabilistic non-numerical calculi such as Norton (2008), and imprecise probabilities such as de Cooman and Miranda (2007). For an overview of approaches see Halpern (2003, Ch. 2) and Dubois (2007).

judgment, which means that we lack any degree of belief whatsoever about any specific possibility or set of possibilities. Then, what I will argue is that in some of these situations of ignorance we should nevertheless also follow the "additivity *of possibilities*", that is, *we should be more confident in the treasure being in* $\{$*chest 5* $\cup$ *chest 6* $\cup$ *chest 7*$\}$ *than in chest 1*.

The type of situations to which this doxastic norm applies are those in which we are in front of a definite number of possibilities but we lack any evidence at all about which is the case. This includes our ignorance of their objective chances, if any. Such situations are called situations of *classical ignorance* (Hansson, 1994). Thus, the doxastic norm here proposed recommends that even if we cannot sum degrees of belief of possibilities, in classical ignorance we should be more confident in any larger group of possibilities.

Conversely, we will also see why this norm does not apply to situations in which we do not know which are the possibilities—that is, in situations of so-called *total ignorance*, stronger than the above classical ignorance. In these situations—which are like being behind a thick, as opposed to a thin, veil of ignorance—we should only suspend judgment. The different recommendations highlight the relevance of the distinction between classical and total ignorance. The result hinges on the rejection of an axiom that has been proposed in the literature of decision theory: Milnor's (1954) AXIOM 8. I argue against its universal applicability in §4.2.1.

A consequence of this norm is that the axiom of non-additivity does not hold. Hence, *in classical ignorance neither additivity nor non-additivity holds.* The agnostic view should thus discard non-additivity in classical ignorance, while keeping non-additivity in total ignorance.

Let us call this doxastic norm (ADDPOSS). It can be schematized in a standard notation thus:

(ADDPOSS): In situations of classical ignorance, for any sets of possibilities A and B, if $|A| > |B|$ then $A \succ B$.

where $|A|$ denotes the number of atomic possibilities in A, and $A \succ B$ means that we are more confident in A than in B. We will assume that the total number of possibilities is finite. We talk of the "additivity of possibilities" in the sense that whilst we cannot sum degrees of belief, summing the atomic possibilities of each A and B allows us to establish a preference for the larger set. We ought to follow (AddPoss), I shall argue, on grounds of *epistemic* rationality. [2]

In what follows, I first spell out the background context (§2) and show that the norm does not fall into well-known objections (§3.1 and §3.2). Then, I discuss the consequences of accepting this norm for the debates between the two aforementioned frameworks for representing our ignorance (§3.3), namely, (1) the agnostic stance portrayed thus far, which we are going to assume (for different defenses and formalisations see e.g. Dubois et al. 1996; Dubois 2007, de Cooman and Miranda 2007, Norton 2003, 2008, 2021, Myrvold 2014, and Benétreau-Dupin 2015), and (2) objective Bayesianism (Urbach and Howson, 1993; Hacking, 2001; Williamson, 2010; Pettigrew, 2016a).[3] More specifically, in §3.3 I derive some consequences of endorsing our norm that lead to the framework of objective Bayesianism, *but* I show that this derivation relies on three additional hidden premises that the agnostic would not accept.

With this in place, I argue why we ought to follow this doxastic norm (§4). To this end, I apply standard decision rules under ignorance, which end up recommending following our norm. This recommendation, though, depends on the rejection of the aforementioned axiom proposed in the literature of decision theory: Milnor's (1954) axiom 8. In a setting

[2]As we will see, the resulting preference order verifies what is known as monotonicity in cardinality. For details on comparative confidence judgments, see Titelbaum (2022, ch. 3). This norm must not be confused with the usual axiom of additivity of comparative orders (in turn different from Kolmogorov's axiom of additivity stated above), which can be stated as follows: for disjoint events A, B, and C, $A \succ B \Leftrightarrow A \cup C \succ B \cup C$ (de Finetti, 1937). Unless stated otherwise, throughout this paper the term 'preference' will thus refer to an *epistemic* preference, i.e. a comparative *confidence*. It will be only in Sect. 6 when we will explore the connections of our epistemic norm with decision-theoretic and pragmatic norms for action.

[3]Other representations of ignorance include Dempster-Shafer's theory of belief functions (Shafer, 1976), or plausibility measures and entrenchment rankings from AGM theory (Rott, 2001; Halpern, 2003) (which do not satisfy additivity, and are only super-additive). For more details, see Genin and Huber (2022, esp.§3).

in which we lay down in columns each possible outcome (see §4.2), this axiom states that identical columns must be merged. Our discussion, then, will serve to clarify the range of applicability of this axiom: it applies in total ignorance, but not in classical. Then, I point out distinctive features of this norm (§5); namely, that it is an epistemic and comparative norm. Having clarified these features, I survey examples in which this norm is applied (§6).

In sum, the territory between the agnostic view and the objective Bayesian view is explored, and the conclusions are that, assuming the agnostic view, (i) in classical ignorance something more than suspension of judgment can be believed on grounds of epistemic rationality: a comparative confidence corresponding to a strict partial order, with the consequence that non-additivity is incorrect; while, on the other hand, (ii) no stronger (more precise) belief structure is warranted—so even in classical ignorance a uniform precise numerical probability distribution is unwarranted. Additionally, (iii) these conclusions rely on the argument that Milnor's Axiom 8 holds only in total ignorance, not in classical ignorance (§4.2), a result which, in turn, stresses the importance of not conflating the settings of classical and total ignorance.[4]

2 The Background: Agnosticism and Non-additivity

To clarify the background context, suppose that an avalanche of rocks divides Indiana's path into two paths such that the possibility space is now divided into two groups of chests: Path A leads to chests [1, 2] and Path B leads to chests [3..10]. Suppose further that Indiana can choose only one of the two paths. (He cannot go back because, say, the floor tiles disappear as he steps out while advancing.) As he is suspending judgment, non-additivity holds: this means that $P(\{1\} \cup \{2\}) = I$ and $P(\bigcup_{i=3}^{10}\{i\}) = I$, where $P(X)$ is the degree of belief that

[4] To further frame the present research, let me underline that there has been an *orthogonal* long discussion between qualitative and quantitative approaches, mostly focused on whether there is an agreement between them. This is an orthogonal debate. The agnostic view, which I assume and explain in this paper and that has been studied in epistemology, decision theory, and philosophy of science, is weaker than any qualitative probability. (Hence, discussions like Kraft et al. 1959 and Scott 1964 are here irrelevant; that is why in Section 3 I explore and show that the lack of agreement between the two frameworks under analysis.) The result is that a specific comparative order is warranted on top of the agnostic view.

an ideal rational agent should have of the treasure being in a chest of set X. Thus, the degree of belief in the chest being down path A and the degree of belief in the chest being down path B is the same doxastic state of agnosticism.

In light of the above, Indiana is agnostic about paths A and B. He is going to follow path A just because it is closer, but at the last minute he decides to call his father Henry for advice. At first, in the same situation of ignorance as his son, Henry agrees with the agnostic stance portrayed thus far:

> Henry: "I have no evidence at all, so I lack any belief about where the treasure is. The treasure might just as well be down path A as down path B. Likewise, I lack any belief about the *chances* of finding it down one path or the other, as I cannot assume that each chest has the *same* chance of hiding the treasure. Hence, it is unwarranted to prefer one path over the other, as this preference would not be based upon any evidence. I would be extracting knowledge from ignorance!"

But then Henry is pressed by his son, thanks to which we are going to discover that there is more to the story:

> Indy: "Ok dad, I agree that we lack any degree of belief of the treasure being down either path, and that we have no clue at all regarding the objective chances, which in turn means that we cannot assume a uniform chance distribution among the chests. *Still*, I am asking you for something *weaker*! Just a *comparative* issue: should I have more confidence in one path over the other?"

> Henry: "I see son. In that respect I recommend you be more confident and therefore choose path B. In this kind of situation *everyone should be more confident in any larger group of possibilities.*"

How could Henry's second answer be correct, in light of his first answer? In the next two sections we will see whether he has committed some well-known mistakes. After that, §4 will show the epistemic rationality of Henry's recommendation.

3 Consequences

Let us now investigate whether there are consequences of endorsing our norm for the debate over representing ignorance between the agnostic and the objective Bayesian.

3.1 Are we Assuming Uniformity?

If the objective chance distribution among the possibilities were uniform, each possibility would have the same weight, so we could sum them and be more confident in the larger group. However, if the unknown objective chance distribution is not uniform but, say, highly peaked in the chests [1,2] of path A, then the sum of possibilities would lead us to a mistaken confidence in path B. Is Henry thus falling prey of the intuitive but unwarranted assumption of a uniform distribution of objective chances?[5]

No. An underlying idea of the Principal Principle is that, if the chances are known, we should match our credences to the objective chances. Yet, Henry does not only does not know the chances, he does not even aim for his credences to approximate the objective chances, as he does not endorse a credence distribution—he just endorses a weak comparative order (cf. §5 below). [6] [7]

Some stop here: under the impossibility of approximating the credences to the objective chances, they suspend judgment. Yet there is, I think, something more that we can, and should, believe.

[5]The lack of justification for a uniform probability distribution has been widely stressed in the foundations of statistical mechanics (see e.g. Albert 2000, Ch. 3.2).

[6]In our example, the objective chance distribution describes the chances of the treasure being inside any of the chests. The physical significance of this distribution can be considered as modeling an unknown underlying stochastic process delivering the treasure to one of the chests. Similarly, this chance distribution can be also considered as modeling a process without genuine objective chances, such as a deterministic process instead of a genuinely stochastic process (corresponding, for instance, to someone putting the treasure in one chest).

[7]The Principal Principle states that, if p is a certain proposition about the outcome of some chancy event and E is our background evidence at t, which must be admissible evidence, then: $Cr(p|Ch(p) = x \wedge E) = x$. (Evidence E is admissible relative to p if it contains no information relevant to whether p will be true, except perhaps information bearing on the chance of p.) See Lewis (1980, 86), Hoefer (2019, Ch. 3).

3.2 Are we Extracting Knowledge from Ignorance?

Another common objection is that, in situations of ignorance, no belief is warranted because we cannot extract knowledge about the world from ignorance. In fact, we cannot; yet this would be a misleading objection to our norm, and now we can see more clearly why. Ignorance, as argues White (2009, 163), puts normative constraints on what our credences should be. This is different from extracting knowledge about the world. In our case this is going to be especially evident, since what I am going to propose is an order of confidence, which cannot and does not aim at matching the objective chances—such an aim would be an unwarranted attempt to extract knowledge from ignorance. It is in this sense of not aiming to form such beliefs about the world that we are not extracting knowledge from ignorance.[8]

3.3 Proofs of agreement

Prima facie, from the point of view of objective Bayesianism, one might think that, along with recent literature that has attempted to derive the principle of indifference (see footnote 12), the present norm also suffices to yield this principle. Yet on the other hand, the advocate of agnosticism could stress an opposite interpretation: yes, in classical ignorance we can say a bit more than mere suspension of judgment, but (AddPoss) is the strongest doxastic commitment we can endorse and, as spelled out in §6, such commitment is weak and of limited applicability—much weaker, the agnostic could say, than the too strong, unwarranted, Bayesian framework.

In what follows, I support the interpretation of the agnostic position. I show that

[8] Some details about our setting may help to prevent the temptation of thinking that the existence of "counterexamples" would undermine the rationality of endorsing our norm. Given our ignorance *we must not make conjectures* about the chances and allow that any chance distribution could be the case. Obviously there are plenty of "counterexamples" in which our recommendation leads to a wrong confidence (i.e., prefer the path where there is no treasure)! Still, the recommendation claims that, given our ignorance regarding the presence or lack of "counterexamples" and therefore acknowledging the obvious fallibility of the norm, it is nevertheless epistemically irrational not to be more confident in the larger group of possibilities. So, strictly speaking, such cases are not counterexamples.

the order recommended by (AddPoss), under one further assumption (the assumption of comparability), can be uniquely represented by a quantitative uniform distribution. However, this is not sufficient to yield Bayesianism. Two further assumptions would be needed. Moreover, none of these three assumptions need be accepted by the agnostic.

First, let us show that, under one further assumption, the unique representation of our resulting comparative order is the uniform distribution. A uniform probability distribution always coincides with our comparative order, in that whenever the size of the sets is larger also the probability of the larger set is higher. Thus, now we have to see that our comparative order can be represented *only* by the uniform distribution. To prove this we would need a condition of comparability like:

$(*)$: If $|A| = |B|$, then $A \sim B$.

As usual in these discussions, $(*)$ sounds intuitive; yet, it does not follow from (AddPoss). (AddPoss) does not constrain what to believe when $|A| = |B|$. The lack of comparability between sets of the same size does not mean that one should have the *same* confidence in them, as I have pointed out in §4.1.

It might also *seem* that one could deduce $(*)$ if the conditional in (AddPoss) were a biconditional, that is, if we assume the other direction:

$(**)$: If $A \succ B$ then $|A| > |B|$.

Assuming $(**)$[9] we have that for any sets of possibilities A and B,

(AddPoss*): $A \succ B \longleftrightarrow |A| > |B|$

[9] A way to justify the assumption of $(**)$ is to maintain that we began our discussion in an epistemic situation of ignorance in which we were suspending judgment; thus, we were assuming that there was no reason that could lead us to any comparative order. Hence, the only reason is (AddPoss); therefore, $A \succ B$ *only if* $|A| > |B|$. This justification, however, seems to conflate our epistemological ignorance of reasons with the objective lack of reasons. Yet, this justification might be acceptable insofar as no other reason has been provided in the literature; that is, in an epistemic situation of suspension of judgment, we assume that there *is no* other reason to endorse such comparative order.

Although it seems so, (AddPoss*) does not imply ($*$). [10] If we had ($*$), we could have concluded that no non-uniform quantitative probability can represent the comparative order determined by (AddPoss*).[11]

Furthermore, *even if* we would have derived such unique representability, still this should not have been misunderstood as a justification of the objective Bayesian probabilistic framework, by which I refer to the 3 Kolmogorov probability axioms and the interpretation of the probabilities in terms of precise degrees of belief. In particular, neither such interpretation nor the axiom of finite additivity would have been justified. We have only proved the unique *representability* of one structure by another, but further epistemic commitments must not be taken for granted unless they are justified, as already Fine (1973, IIID) stressed and Norton (2003, 2007, 2008, 2010, Forthcoming) has been urging (see also the explicit diagnosis by Titelbaum's (2022, Ch. 2.4 and Ch. 10.3.3) book on Bayesianism concerning the difficulties in justifying the axiom of finite additivity). We must not misinterpret a precise numerical uniform distribution as giving us too precise—unwarranted—information: changing the non-numerical symbol 'I' with numbers is not necessarily more than a *mere* re-labeling. If one considers that the uniform numerical probability provides novel information, then one could raise the question: is this new

[10] One might be tempted to conclude ($*$) due to the following invalid proof by reductio. By switching the names of the variables in (AddPoss*) we have the logically equivalent version with the signs changed (AddPoss**): $A \prec B \longleftrightarrow |A| < |B|$. Let us show that it is impossible that $(*) : |A| = |B| \to A \sim B$ is false and both (AddPoss*) and (AddPoss**) true. Let A and B be such that $|A| = |B|$ is true and $A \sim B$ false. Then, ($*$) is false, $|A| > |B|$ and $|A| < |B|$ are false, and either $A \succ B$ or $A \prec B$ is true. But then either (AddPoss*) or (AddPoss**) is false. Therefore ($*$). ■ (Note that if we use (AddPoss) instead of the stronger (AddPoss*), the proof does not follow, since ($*$) would be false and the premises (AddPoss) and the analogous of (AddPoss**) vacuously true.) This proof is, however, incorrect. Again due to an intuitive but wrong disregard of the agnostic position, the proof is *wrongly* supposing that if $A \sim B$, then it must be the case that either $A \succ B$ or $A \prec B$. This is wrong, for in fact one can suspend judgment and have no confidence whatsoever about the relation between A and B.

[11] Proof. Consider a generic non-uniform quantitative probability P. Then, for some atomic possibilities $a, b \in \Omega$ (where Ω is the finite possibility space), $P(a) > P(b)$. Since a and b are atomic possibilities, $|\{a\}| = |\{b\}|$. By ($*$), we have that $a \sim b$. Then, it is not the case that $a \succ b$. Therefore, it is false that $P(a) > P(b)$ iff $a \succ b$. Hence, no non-uniform P can represent the order determined by (AddPoss*). ■

information warranted by the scenario being modeled? The answer, in our scenario, is no: using numbers and the same number for each atomic possibility justifies neither the interpretation that the numbers represent the precise strength (or 'degree') of an agent's belief nor their additivity—without a justification, these commitments will not be accepted by an advocate of suspension of judgment. This can be clearly seen by considering the examples surveyed in Section 6, in which we are endorsing (AddPoss) but a precise and finitely additive uniform distribution of degrees of belief is not warranted. In particular, in the examples of Clifford's boat (§6.2) and in scientific and philosophical inquiry (§6.4), we explicitly saw that we were not warranted to extract precise information from the comparative order, an information which a fortiori we cannot sum.

4 Justifying the Norm by Minimising the Risk of Inaccuracy

Could Indiana, *in addition to* suspending judgment about where the treasure is, hold another epistemically rational belief? While the most cautious doxastic state will be shown to be suspension of judgment, another rational doxastic state following cautious rules will be shown to be *a comparative order of preference* between groups of possibilities, ordered according to their size. This order is a strict partial order on the power set of possibilities, i.e. of chests. In this section the epistemic rationality of this comparative belief is justified.

4.1 You can't lose by picking one extra chest

First of all, we are going to consider the *accuracy* of our beliefs. It is rational for Indy, in his aim to believe the best guess regarding the treasure's location, to follow the doxastic norm of *minimising the risk of inaccuracy* of one's belief in a proposition. The notion of accuracy is not a pragmatic but an epistemic and alethic value of a doxastic state: it indicates how much an agent's doxastic state approximates to truth.[12]

[12] Similar but more sophisticated arguments based on the notion of accuracy have been proposed to defend a stronger claim: the assigning of equiprobability in situations of ignorance. Pettigrew (2016a,b) directly appeals to accuracy, while Williamson (2010, §3.4) and Landes and Williamson (2013) appeal to a pragmatic caution understood as the minimisation of loss, although Williamson (2018) adapts the argument

Before the more general argument in §4.2, which is made within the formal framework of decision theory, consider the following argument. Think of a single chest: imagine that for some reason you are allowed to add a single chest from the other group to the ones you have. Doing so *only adds an opportunity* of getting the treasure. In epistemic rather than decision-theoretic terms, believing that the treasure is in a group with an additional chest only adds an opportunity of believing in the right chest. Likewise, picking an extra chest *cannot* decrease our opportunities of being right. When we believe that the treasure is in a group with an additional chest, our belief does not necessarily approach the truth—it might approach it or it might remain the same. Yet this belief and the previous belief without the additional chest cannot be valued equally: the opportunities of having an accurate belief have increased, that is, *our risk of inaccuracy has decreased.*[13] Thus, the more chests we pick, the less we risk holding a false belief. And this extra opportunity comes for free. Hence, in comparing groups in which one is a strict subset of the other, we should have more confidence in the larger group.

This first argument, defending the additivity of possibilities only for strict subsets, would already threaten the consistency of the axiom of non-additivity, *if* we accept transitivity and comparability between sets. For the agnostic would, for instance, maintain both that $\{1\} \cup \{2\} \sim \bigcup_{i=3}^{10}\{i\}$ and that $\{1\} \cup \{2\} \cup \{3\} \sim \bigcup_{i=4}^{10}\{i\}$. But then, by the additivity of strict subsets, we have (a) $\{1\} \cup \{2\} \cup \{3\} \succ \{1\} \cup \{2\}$ and (b) $\bigcup_{i=3}^{10}\{i\} \succ \bigcup_{i=4}^{10}\{i\}$. Therefore, we have $\bigcup_{i=4}^{10}\{i\} \sim \{1\} \cup \{2\} \cup \{3\} \succ \{1\} \cup \{2\} \sim \bigcup_{i=3}^{10}\{i\}$. But if we accept transitivity, $\bigcup_{i=4}^{10}\{i\} \succ \bigcup_{i=3}^{10}\{i\}$ is inconsistent with (b).

However, the fact that the agnostic has no degree of confidence whatsoever about A and no degree of confidence whatsoever about B must *not* be misunderstood as that the agnostic having the *same confidence* in A as in B. That is, the agnostic refuses to

to a purely non-pragmatic epistemic version appealing to accuracy. I do not assume these stronger results, which would imply our weaker claim (for a numerical uniform distribution implies that the larger set is more likely). See also Joyce's (1998) argument for probabilism from accuracy, and Pettigrew (2016a) and Titelbaum (2022, ch.10) for a clear treatment of the notion of accuracy.

[13] As is usually understood in the literature, 'risk' just means the possibility that some undesirable event may occur (Hansson, 2018, §1), in this case related to the inaccuracy of our beliefs.

	w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8	w_9	w_{10}
A	+k	+k	-k	-k	-k	-k	-k	-k	-k	-k
B	-k	-k	+k	+k	+k	+k	+k	+k	+k	+k
Suspension	0	0	0	0	0	0	0	0	0	0

Table 1: Accuracy Matrix of three beliefs that Indiana could have, for each of the 10 possibilities w_i, where each w_i corresponds to the treasure being in chest$_i$. The value '+1' would correspond to maximum belief accuracy (that is, total correspondence with an omniscient agent's belief, that is, true belief). The value '+k' corresponds to the value of a certain degree of accuracy but not the maximum. The specific value doesn't matter: all is needed is that true belief is better than suspension of judgment, which is in turn better than false belief.[14]

endorse comparative statements '$A \sim B$' concerning any sets A, B. One might be tempted to draw such mistaken conclusions due to a potential misinterpretation of the formalism presented above (from Norton, 2008). Given that one writes down things like '$\{2\} \sim I$' and '$\{1\} \cup \{2\} \sim I$', and assuming some sort of transitivity and commutativity, it seems natural to conclude that $\{2\} \sim \{1\} \cup \{2\}$. However, that would be fallacious, as all the steps in this inference—commutativity, transitivity, and overly strong readings of statements like '$\{2\} \sim I$'—can be disputed by advocates of suspending judgment.

4.2 Decision theory applied

The previous argument has compared two sets of possibilities in which one is a strict subset of the other. The same conclusion—the preference for the larger set—can be independently obtained within the framework of decision theory for any sets of possibilities. The means are the same: minimising the risk of inaccuracy. Table 1 describes our case study, summarizing the accuracy of three comparative beliefs in the rows ('be more confident that the treasure is down path A', 'be more confident that the treasure is down path B', and 'suspend judgment as to where it is'), according to what turns out to be the case in the columns.

[14]If you say $A \succ B$ and it is in A, then you are accurate, but not maximally accurate ('+1') because your claim was weak; so your accuracy scores some positive value '+k', $0 < k < 1$. If, instead, you say

To the extent that the aim is true belief, and avoiding false beliefs takes priority over holding true beliefs, it is reasonable to follow the most cautious rules. (Only in the epilogue, in §6.5, will I relax this cautious attitude while still aiming at truth.) We then have to avoid rules like MAXIMAX, and follow DOMINANCE, MINIMAX, MINIMAX-REGRET, LEXIMIN, and HURWICZ (with a low, i.e. cautious, value for its parameter α).[15]

We stipulate a value of 0 to the accuracy of suspending judgment. Given that, one can verify that all rules determine suspension of judgment as the most cautious doxastic attitude to endorse, as I said previously. But then, DOMINANCE, MINIMAX, and MINIMAX-REGRET yield no preference for A or B, disagreeing with my main claim that B is to be preferred to A. We can verify this lack of preference by noting that these rules focus only on the worst values of the matrix; hence, this allows merging identical columns as in table 2 (a move discussed below, known as 'merger of states' or 'column duplication'); hence, this leads to a symmetrical situation with no preference between beliefs A and B.[16]

$A \succ B$ and it is in B, then you are wrong, and so inaccurate to a certain degree. The score would not be maximally inaccurate ('-1') for the same reason (your claim was merely a weak comparative statement, not an assignment of a precise credence distribution). It would instead be some negative accuracy value: let us say '-k'. For more on this see, along the same lines, Konek (2015, p.2 and Sections 3,4). In any case, potential variations in the specific values would not affect the argument below, since the decision rules would yield the same results.

[15] MAXIMAX is a rule associated not with caution but with taking risks, since what it recommends is to prefer the option with the greatest utility, even if its worst possible outcome is worse than those of its alternatives. In contrast, the other rules are cautious in that they prioritize avoiding the option with the worst utility. We will discuss these rules later, when I apply them to our scenario. In any case, for a discussion of whether these are the most cautious rules, see Williamson (2007) for synchronic caution and Radzvilas et al. (forthcoming) for diachronic caution.

[16] In more detail, these rules determine the following (in both tables): there is neither strong nor weak DOMINANCE, for neither A nor B are in all worlds as good or better than the other. MINIMAX makes no choice either, for it would choose the path with the best worst-case inaccuracy, but that is '-1' for both A and B. MINIMAX-REGRET makes no choice either, for the worst-case regret is the same (the worst-case regret of A is +1- -1=+2; *idem* for B). Finally, HURWICZ takes into account the worst and the best values of each option, conferring a relative importance to each through a parameter α: $A \succ B$ if and only if $\alpha \cdot max(A) + (1-\alpha) \cdot min(A) > \alpha \cdot max(B) + (1-\alpha) \cdot min(B)$. Also HURWICZ recommends equally both A and B, no matter the value of α, for in the formula $\alpha V_{MAX} + (1-\alpha)V_{min}$ the maximum and minimum accuracy values V_{MAX} and V_{min} are the same for A and B. We will now see that the tie is broken when we

	$w_1 \cup w_2$	$w_3 \cup ... \cup w_{10}$
A	+k	-k
B	-k	+k
Suspension	0	0

Table 2: Accuracy Matrix after merging identical states (columns). As before, all cautious rules recommend suspension of judgment; and due to the symmetry no rule determines a preference of B over A or vice versa.

However, a preference between A and B is obtained once we apply the more sophisticated Leximin rule. This is because under classical ignorance Leximin has to be applied, as I argue below, only to table 1—in other words, merging identical states must not always be allowed.

Leximin is the rule that holds that if the worst-case inaccuracies between options are the same, one should choose the option such that its second worst-case inaccuracy is minimum; if this is not a tie-breaker, then the third worst-case inaccuracies should be compared, and so on:

Leximin:
$A \succ B$ if and only if there is some positive integer n such that
$min^n(A) < min^n(B)$ and $min^m(A) = min^m(B)$ for all $m < n$

where $A \succ B$ means that we should be more confident in A than in B, $min^1(A)$ is the worst inaccuracy of belief A, $min^2(A)$ is its second worst inaccuracy and $min^n(A)$ is its n^{th} worst inaccuracy (cf. Peterson 2009, 45).

Hence we have the result that, of the rows '$A \succ B$' and '$B \succ A$', Leximin recommends $B \succ A$. That is, also from the point of view of epistemic rationality, not all paths are on a par.

One could stick to the other rules previously considered and thereby keep the indifference between A and B—but this does not seem *the most* reasonable choice. For Leximin can

recur to the remaining rule Leximin.

be seen as a more sophisticated version of the other rules, insofar as it delivers deeper, more fine-grained, prescription about the options under scrutiny. So it is hard to see why one would put aside Leximin and stick to the tie determined by any other rule. We have seen that there *is* a reason to avoid the tie, so it seems unreasonable to sweep it under the rug. Imagine for instance that one insists on the preferability of Minimax for some reason. Not only does no such reason actually exist; such rules are neither radically different nor incommensurable, so anyone can compare them and thus see that it would not be *the most* reasonable choice to favour the more simplistic over the more fine-grained rule. In fact, ceteris paribus, one would have to pretend that the additional information provided by Leximin is not worth taking into consideration.

Then, as mentioned before, in table 1 suspension of judgment correctly turns out to be the most cautious doxastic attitude (as it always is in situations of uncertainty); and as we will clearly see from the examples of Section 6, the comparative order says nothing about many propositions about which one should still suspend judgment—hence the compatibility between (AddPoss) and the agnostic view.

4.2.1 Objection: merge identical possibilities

One could object that Leximin may be applied only to table 2: since the chances are unknown, why not conflate identical possibilities? Let me illustrate with a further example why, while this move is often either mandatory or not mandatory but innocuous, it sometimes misrepresents the case under consideration by leaving out relevant information.

In other words, I am going to argue that the validity of Milnor's (1954) axiom 8, which states that identical columns must be merged, does not always hold.[17] That is, this axiom—not one of the indispensable axioms, according to Milnor himself and Binmore (2008, 158)—does not hold for every kind of ignorance, but only for the most extreme total ignorance, when we do not even know which are the real possibilities. A possibility is real if

[17]More specifically, the axiom states: "`8. Column duplication. The ordering is not changed if a new column, identical with some old column, is adjoined to the matrix. (Thus we are only interested in what states of Nature are possible, and not in how often each state may have been counted in the formation of the matrix.)`" (Milnor, 1954, 52).

and only if there is an associated nonzero objective chance. That is, it is a possible outcome of the laws of Nature—in current terminology, it is physically possible.[18] Cf. Milnor 1954, 56, who stresses that on some occasions "it is desirable to preserve axiom 8", especially when "there is no clear and natural separation of the possible states of Nature into a finite number of finite alternatives." See also Luce and Raiffa 1957, ch. 13.3 to 13.5, for discussion of this and other axioms.

The example comes from Ellsworth (1978), who studied Rawls' veil of ignorance, and in so doing he studied utility tables essentially identical to ours. He presented a table like table 1 (*ibidem*, 34) in which two rows, A and B, represent two societies, each column represents a social class, and the wealth of each social class is the utility quantified in the cells. The famous question posed by Rawls is which society is better for the agent to choose, given her ignorance of to which social class w_i she will belong.

One might be tempted to answer that since the chances of each w_i are unknown, we can merge identical columns as in table 2 and therefore we can have a preference neither for society A nor B. However, this lack of preference between societies relies on our formal representation of the veil of ignorance: a *thick* veil of ignorance—i.e., *total* ignorance—requires that we conflate identical columns (some of which might have zero chance). But a *thin* veil of ignorance—i.e., *classical* ignorance—considers all possibilities as real possibilities (that is, having a nonzero chance of being the case). And Rawls' scenario, according to Ellsworth (1978), is one of classical ignorance—just like Indiana's scenario.

More specifically, in total ignorance it is indeed mandatory to conflate identical possibilities, because we ignore which are the real possibilities. In classical ignorance, instead, we cannot always conflate two *real* possibilities into a single possibility. In classical ignorance it also *seems* appropriate to conflate them as regards their chances, insofar as the conflation preempts thinking that the chance of the sum of, say, a hundred possibilities is larger than the chance of one possibility. However, in so doing we also remove information

[18] Not all possibilities are real, they can be just merely conceptual (or logical): they *cannot* occur; that is, they are *not physically possible*; in other words, they are not an outcome of an objective chance distribution. Now, we do not know a priori what is physically possible, and in many scenarios we never know which is the space of real, physical, possibilities.

regarding *which* are the real possibilities. And while chance information is irrelevant in the application of LEXIMIN, *which* are the real possibilities constitutes relevant information: the indices m and n in LEXIMIN's formula range over *all* those possibilities. Thus, the conflation of possibilities is often innocuous (as when applying MINIMAX) or even mandatory (as when ignorant of the real possibilities), but it is not always correct: at least, not when applying LEXIMIN.

Following Peterson (2009, 49), a similar argument can be constructed appealing to a sophisticated version of HURWICZ. The standard version of this rule takes into account the worst and the best values of each option, conferring a relative importance to each through a parameter α (for its definition see footnote 16). In regard to this rule, Peterson (2009, 49) raises the objection I raised before: "The obvious objection is, of course, that it seems irrational to totally ignore the outcomes in between. Should one really ignore all the non-minimal and non-maximal outcomes?" He then also stresses that one could appeal to the unknown chances to conflate the columns; however, "the basic objection remains. Why is it rational to focus exclusively on the best and the worst possible outcomes? What about the other, non-extreme outcomes?" Then, he responds with a proposal that assigns α-values to *all* outcomes, not only to the best and the worst. The α-values are not to be interpreted as probabilities, but as the importance that should be given to each outcome. For more details see Peterson (2009, 49).

4.3 Arbitrary partitions

The discussion regarding tables 1 and 2 reminds us of the partition problem, the problem that ends up yielding inconsistent recommendations depending on how we carve up the possibilities. Even if we are not applying the principle of indifference, we are not immune to the threat of the partition problem. Yet, in the kinds of scenarios considered, we do not face the problem because the scenarios are endowed with a set of real possibilities. This allows us to avoid the problem since we can, *ex hypothesi*, identify *the* partition. The harder case of total ignorance is, on the other hand, inevitably unable to provide a best single partition.

4.4 Summing up

In the end, after assessing all the main cautious decision rules under classical ignorance (not only Leximin), the minimum risk of worst-case disutility is obtained by choosing the society where 8 social classes w_i are rich and 2 are poor, not the other way around—i.e., preferring society B. Likewise, the same rules, with Leximin breaking the tie imposed by the others, determine that the minimum risk of worst-case inaccuracy is obtained by being more confident that the treasure is down the path with 8 chests, not down the path with 2 chests—i.e., preferring path B.

Our rationale has been to survey all the cautious decision rules under ignorance, whereby we found that all the rules agreed that the most cautious option is always to suspend judgment, and then one of these rules added some further information, namely that A and B do not weigh the same. Most rules yield a tie between the options and one of them, Leximin, goes beyond the tie and tips the balance. This led us to conclude that *not all paths are on a par*.

This argument in §4.2 is independent from and adds up to the argument limited to strict subsets in §4.1 according to which, by assessing the risks of adding a single chest, it was concluded that we can only minimise the risk of inaccuracy by adding more chests. It is worth noting that in the case of strict subsets, the agnostic is *not* permitted to remain agnostic about whether $A \succ B$. For there is no uncertainty, no risk, about whether the agent has a wrong preference. However, the agent can still be agnostic about anything else. This contrasts with the general case, in which we have seen that one can still remain agnostic about being more confident in A than in B, if one aims to maximise caution. Accordingly, in Section 6 we will survey cases concerning when to remain agnostic rather than endorsing the comparative order.

Summing up, in classical ignorance we should suspend judgment but we are also justified, under the criteria of cautious (and non-cautious) epistemic rationality, to endorse a strict partial order of comparative confidence among sets of possibilities: the larger the set, the better, for this reduces the risk of holding inaccurate beliefs. To be more confident in a smaller set would be to take an unnecessary risk of false belief, on pain of epistemic irrationality.

5 An Epistemic, Comparative Norm

Having justified the norm (AddPoss), let us survey some of its main features to better understand its potential applicability. In light of these features, in Section 6 we will survey some examples to ascertain in which type of cases the norm is useful and in which cases it is not.

5.1 An epistemic norm

As it has been defined (and as it has been justified and as the following examples will make evident), (AddPoss) is an epistemic norm. It is not the norm of minimising the risk of moving away from some pragmatic value, such as survival. Rather, it is the norm of minimising the risk of moving away from truth. Thus it is not only that Indiana the adventurer should better choose path B to save his life; it is also that Indiana the scientist should be more confident that B, rather than A, contains the "true" chest. Thus, the norm is compatible with the widely held monist view called 'veritism', according to which truth is the unique aim of belief.

5.2 A non-evidentialist norm

At the same time, it might be thought that the norm sides with non-evidentialism, a position in the field called 'ethics of belief'. This area proposes norms—mainly of an epistemic, pragmatic, or moral sort—that recommend what we should believe (see Chignell 2016 for an overview). Evidentialism is the (reasonable and widely-held) recommendation not to believe anything on insufficient evidence. A moderate version of evidentialism allows that in some contexts it is permissible to have beliefs based on insufficient evidence. Yet the doxastic norm proposed here *seems* to side with non-evidentialism, which claims not just that it is *permissible*, but that in some contexts we *should* hold certain beliefs on insufficient evidence.

To illustrate the non-evidentialist position with a paradigmatic (and disputed) example, we can cite William James, who defended a non-evidentialist norm. He warned of "the risk of losing truth" when being too cautious (James, 1979, 25). This risk is in direct tension with the risk of being in error. While a cautious attitude prioritizes avoiding error, James argued that this priority is disputable, and instead opted to prioritize not losing truth. In the epilogue in §6.5 I assess how my norm would combine with James's alternative priority of not losing truth.[19]

That (AddPoss) is a non-evidentialist norm is, however, disputable. For it could alternatively be said that our scenario sides with evidentialism, in the sense that the conclusions we draw are based on a certain kind of evidence (even if tenuous); namely, the evidence that the possibilities we are considering are (as explained in §4.3) the real possibilities. For a proposal along these lines see Williamson (2010, ch. 9, especially pp. 156-7).

In any case, besides the size of the space of possibilities, we have no evidence about the chances, so in this sense I think the norm can be considered as non-evidentialist.

5.3 A comparative norm: weak, imprecise, and fallible

It is worth noting that (AddPoss) is a weak, imprecise, and fallible norm—usual characteristics of comparative orders of confidence. It is weak in the sense that being a comparative order of confidence, it does not *quantify* how likely it is that A or B is the case, as quantitative (i.e., numerical) probabilities do. It is imprecise in that being a comparative order of confidence, it does not tell us *how much* more confident one should be in B than in A. And of course it is fallible, in that it can be wrong (i.e. recommend us what is not the case).

Unlike the *total* orders implied by a precise uniform probability distribution, (AddPoss) imposes a *strict partial* order, as visualized in Figure 1. Notably, since (AddPoss) is silent

[19]Some differences between James's norm and mine: James defended belief on insufficient evidence without specifying which belief we could hold (albeit having in mind religious belief). Here, instead, I defend holding a *specific* belief on insufficient evidence, and one that is much less ambitious than religious belief (as I clarify in §5.3): a partial order of confidence among sets of possibilities, ordered according to their size.

about the elements of the same size, they lack horizontal lines connecting them, meaning that they are not comparable. That is, we are not warranted to have the same confidence in them; we should keep suspending judgment about them and about their relation.

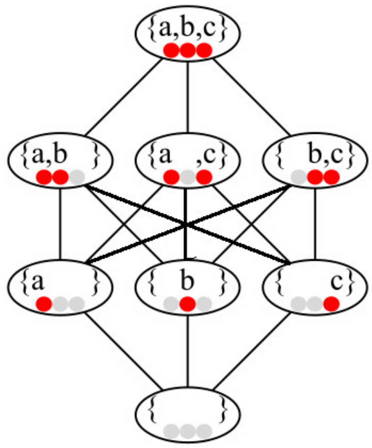

Figure 1: The strict partial order that (AddPoss) recommends to hold for 3 possibilities. The lack of horizontal lines among sets of same size indicates that such sets are not comparable, since (AddPoss) says nothing about them.

In light of these characteristics, the norm proposed sometimes will be relevant in the beliefs we should endorse about a certain issue and sometimes it won't; and similarly, sometimes it will be relevant in deciding to follow some actions (in combination with pragmatic norms for action) and sometimes it won't. Let us survey examples of different sorts.

6 Applying the norm

6.1 Forced to choose

Our example of Indiana Jones is one of those cases in which the agent is forced to choose an option (remember that the floor tiles behind him have disappeared), which is usual in many contexts of decision theory, game theory, artificial intelligence, etc. In such cases, the most cautious doxastic state—suspending judgment as to the location of the treasure—is excluded.[20] But then, do all paths represent an equally rational choice? What has been argued in §4 is that it is epistemically rational to be more confident in the path that minimises the risk of inaccuracy and thus in the path with more chests. The upshot is that in cases in which one is forced to choose, one should not choose arbitrarily—not all paths are on a par.

6.2 Not forced to choose; unwilling to risk

What about contexts in which one is not forced to choose? That is, what about contexts in which one can still endorse the most cautious state, that is, suspend judgment? Consider the ship captain mentioned by Clifford (1877, 1), whose decision to set sail involved risking hundreds of passengers' lives (because the ship had not been properly supervised). He used this example (which actually happened to him—the ship sank but Clifford survived) to expose the defects of taking decisions on insufficient evidence—against non-evidentialism, that is. (AddPoss) sides with the comprehensibly heretical non-evidentialist view, so let us see whether some decision-theoretic norm might recur to (AddPoss) to recommend the (at least apparently) undesirable choice of setting sail. (As I elaborate below, (AddPoss) does not recommend any *action* at all as it is just an *epistemic* norm, but other *decision norms* following (AddPoss) might do.)

Suppose that most of a finite list of possible courses of action are safe—that is, correspond to the ship not sinking. Now, according to (AddPoss) we should be more confident in any larger group of possibilities being the case—so, we should be more confident that the ship will not sink. While according to (AddPoss) it is correct to hold this belief, does this belief imply that the captain should set sail?

Not at all; the final *decision* is in the end the result of some *decision norm* (which presumably will take into account the risk one is willing to take), whereas our norm is just an epistemic norm. Whatever it is the decision norm the agent follows, it will be

[20] In Savage's sense, this kind of situations can be described as those in which the agents do not have the so-called null-act in their current set of possible actions.

informed by the confidence order suggested by (AddPoss), but then the outlined features of our confidence order illustrate why the captain should not set sail (in spite of being more confident that the ship will not sink): we know neither how much more confident we are than in the alternative (due to the imprecision of comparative orders), nor how likely it is that the alternative shall occur (due to the weakness of comparative orders). And here the priority is to maximise caution given that hundreds of lives are at stake (in other words, the extremely negative utility of the ship sinking is far worse than the positive utility of the ship completing the voyage); hence, every sensible decision norm should urge the captain to not set sail.

6.3 Not forced to choose; willing to risk

Yet we do not always need to maximise the caution of our actions. We could choose the larger group—i.e., it would be at least permissible—in any situation in which we are willing to take the risks, such as when nothing like hundreds of lives is at stake. Should you go to a party dressed in red believing that your friend will not dress in red, which is one of the five colours he always uses? According to (AddPoss), we can be more confident that your friend will not dress in red (4 colours vs. 1). Still, lacking any further information, to dress in red is a risky action in that you might go to the party dressed like your friend (something supposedly bad). However, this can be considered a risk that you can take (in other words, the negative utility of dressing the same colour is not remotely as negative as the utility of the ship sinking).

6.4 Scientific and philosophical inquiry: neither forced to choose nor willing to risk

Now, consider the demanding contexts of scientific and philosophical inquiry: our aim is true belief; we have to maximise our caution; and we are not forced to choose. They are paradigmatic contexts for evidentialism (especially, science: to proceed based on evidence is a defining trait of scientific methodology). Imagine that Indiana the scientist returns from his adventure without opening any chest and writes a scientific paper claiming (among other things) that the treasure is more likely to be down path B. Could he establish such a conclusion?

No, if he means by 'likely' the objective chances of the treasure being in B: as discussed in §3.1, he cannot be so ambitious as to guess about totally unknown chances, so talk of 'likeliness' here can be maintained at most only if he refers to our credences. No, if he purports to assign a specific degree of belief to the proposition that the treasure is down path B: this claim refers to credences, but it is still giving information more precise than we are warranted to justify, as explained §5.3. No, if he purports to give a specific value to how much more confident he is in B than in A: again, this is information more precise than we are warranted to justify. Nevertheless, as long as what he claims *is just* that we should be more confident in path B than in A, in the weak sense specified throughout this section, then yes, he can establish such a conclusion.

Why can Indiana claim that we should be more confident that the treasure is down path B? Now it is worth insisting, especially in scientific and philosophical inquiry, that it is obviously safer to remain agnostic. However, this does not undermine Indiana's claim. It is compatible. For we have a reason, a non-evidential epistemic reason, to believe that $B \succ A$, namely the doxastic norm (AddPoss), justified in section 4. In sum, the preference for path B, properly understood, is also correct in an epistemic context such as a scientific article.

6.5 Epilogue – not forced to choose, but willing to believe

The defense of the comprehensibly heretical combination of epistemic and non-evidentialist doxastic states can be pushed further: our weak, imprecise, and fallible comparative confidence and William James' "will to believe" can mutually strengthen one another.

James' "will to believe" is an attempt to legitimize the endorsement of a belief in theoretical contexts on insufficient evidence. He appealed to "the risk of losing truth" (James, 1979, 25) when we suspend judgment, to argue that the methodological principle of avoiding error—which recommends suspension, at the risk of losing the truth— is not *objectively better* than the methodological principle of seeking the truth—which recommends believing, at the risk of believing falsely. (Lacking an objectively better option, James personally prioritized the latter.)

To illustrate how this mutual strengthening could work, imagine that Henry Jones becomes aware of his imminent death and thus wishes to endorse a belief about a breath-taking philosophical hypothesis that has fascinated him all his life. Lacking any evidence, though, it may seem hard to accept James' suggestion. For it is arbitrary to take a stance on the hypothesis. It seems that Henry should still suspend judgment—even if he is going to die and he really wishes to believe.

However, if the hypothesis under scrutiny holds in a certain subset of the possibility space whose cardinality is bigger (or smaller) than the rest of the possibility space, then he has at his disposal our non-evidential norm (AddPoss) to allow him to take a stance—his belief would not be arbitrary anymore. Henry would be justified on grounds of epistemic rationality to endorse a belief—albeit just a weak, imprecise, fallible comparative confidence—without having to apply the principle of indifference.

In short, Henry has avoided the cautious suspension of judgment by means of James's will to belief, and then has endorsed not whatever belief he wills, but that which is recommended by our epistemic norm.

7 Conclusion

The doxastic norm (AddPoss) has been proposed (§1) and defended (§3 to §4): this norm recommends holding, in situations of classical ignorance, a strict partial order of confidence on the power set of possibilities. It has been shown why this norm applies in contexts of classical ignorance and why it does not apply in contexts of total ignorance (§4). Its defense applies standard decision rules under ignorance, where the utility value is understood in terms of the epistemic notion of accuracy (§4.2). This defense involved the discussion of the scope of one of the axioms of decision theory under ignorance, Milnor's (1954) axiom 8, arguing that such an axiom is not valid under classical ignorance.

Suspending judgment is always the most cautious doxastic state. This, however, does not mean that the comparative order of confidence recommended by (AddPoss) is epistemically irrational (§4). Accordingly, we have seen different situations of classical ignorance in which (AddPoss) is and is not relevant for choosing a certain action: on the one hand, it is relevant for deciding the action when we are forced to choose (§6.1: Indiana's case), and when we are not forced to choose but we are willing to accept risk (§6.3: dressing for a party); on the other hand, it is irrelevant for deciding the action when we are not forced to choose and not willing to risk (§6.2: Clifford's boat captain). Still, even in the last type of context, in which we are neither forced to choose nor willing to risk—which includes contexts of scientific and philosophical inquiry—it has been argued (§6.4 and §6.5) that, even if the most cautious doxastic state is suspension of judgment, when we know which are the possibilities it is epistemically rational to have greater confidence in larger groups of possibilities. Moreover, if the options are such that one is a strict subset of the other (as in §4.1), it is not even permissible to keep suspending judgment over the comparative order. That is, no matter the context, one must be more confident in the larger superset—even the most judgment-averse agnostic. Accordingly, take for instance the aforementioned approach of imprecise probabilities, which provides a formal framework for the agnostic view. Such an approach does not take into account the cardinality of atoms; thus, while I think it is an adequate framework for the context of total ignorance, in this context it is "unnecessarily imprecise".

Finally, does (AddPoss) suffice to lead us from the agnostic stance to the framework of objective Bayesianism? No. In §3.3 it was shown that this result would also require us to assume three further conditions that the agnostic does not have to accept, i.e. a comparability condition, the axiom of finite additivity, and interpreting the probabilities as precise degrees of belief.[21]

[21] **Acknowledgements:** For their helpful comments, I would like to thank Brice Bantegnie, Carl Hoefer, John Horden, Kine Josefine Aurland-Bredesen, Wayne Myrvold, Elías Okon, Michael Pockley, Nils-Hennes Stear, Carlos Romero, Sven Rosenkranz, Marta Sznaider, David Teira, Alessandro Torza, an anonymous reviewer of 'Philosophy of science', and audiences at the ALFAN conference, the 'Methodology of philosophy' seminar at UNAM, and the LOGOS seminar.

Funding: This work was supported by the grant 'Formal Epistemology – the Future Synthesis', in the

References

Albert, David Z. *Time and Chance.* Harvard University Press, 2000.

Benétreau-Dupin, Yann. "The Bayesian Who Knew Too Much." *Synthese* 192, 5 (2015): 1527–1542.

Binmore, Ken. *Rational Decisions.* Princeton University Press, 2008.

Chignell, Andrew. "The Ethics of Belief." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Stanford University, 2016. Winter 2016 edition.

Clifford, William K. *The ethics of belief and other essays.* Prometheus Books, 1877, 1999 edition.

de Cooman, G., and E. Miranda. "Symmetry of models versus models of symmetry." In *Probability and Inference: Essays in Honour of Henry E. Kyburg*, edited by W. Harper, and G. Wheeler. College Publications, 2007, 67–149.

Dubois, D., H. Prade, and P. Smets. "Representing Partial Ignorance." *Trans. Sys. Man Cyber. Part A* 26, 3 (1996): 361–377.

Dubois, Didier. "Uncertainty theories: a unified view." In *IEEE Cybernetic Systems Conference, Dublin, Ireland.* 2007, 4–9.

Ellsworth, Leon. "Decision-Theoretic Analysis of Rawls' Original Position." In *Foundations and Applications of Decision Theory*, edited by A. Hooker, J. J. Leach, and E. F. McClennen. D. Reidel, 1978, 29–45.

Empiricus, Sextus. *Outlines of Pyrrhonism.* Harvard University Press, I c. A.C., 1994 edition.

Filomeno, Aldo. "The Propagation of Suspension of Judgment. Or, Should We Confer Any Weight to Crucial Objections the Truth-Value of Which We Are Ignorant?" *Erkenntnis* 1–22.

Fine, Terrence. *Theories of Probability: An Examination of Foundations.* Academic Press, 1973.

de Finetti, Bruno. "La Prévision: Ses Lois Logiques, Ses Sources Subjectives." *Annales de l'Institut Henri Poincaré* 17 (1937): 1–68.

Friedman, Jane. "Rational Agnosticism and Degrees of Belief." *Oxford Studies in Epistemology* 4 (2013a): 57.

———. "Suspended Judgment." *Philosophical Studies* 162, 2 (2013b): 165–181.

———. "Why Suspend Judging?" *Noûs* 50, 3.

Genin, Konstantin, and Franz Huber. "Formal Representations of Belief." In *The Stanford Encyclopedia of Philosophy.* 2022. Fall 2022 edition.

Hacking, Ian. *An Introduction to Probability and Inductive Logic.* Cambridge University Press, 2001.

Halpern, Joseph Y. *Reasoning About Uncertainty.* MIT Press, 2003.

Hansson, Sven Ove. "Decision theory." *A brief introduction. Department of Philosophy and the History of Technology. Royal Institute of Technology. Stockholm* .

framework of the program Praemium Academicum realized at the Institute of Philosophy of the Czech Academy of Sciences, and by the DI-consolidado 039.310/2022 of the Pontificia Universidad Católica de Valparaíso.

———. “Risk.” In *The Stanford Encyclopedia of Philosophy*, edited by Ed Zalta. Metaphysics Research Lab, Stanford University, 2018. Fall 2018 edition.

Hoefer, Carl. *Chance in the World: a Skeptic’s Guide to Objective Chance.* Oxford University Press, 2019.

James, William. *The Will to Believe and Other Essays in Popular Philosophy.* Harvard University Press, 1979.

Joyce, James M. “A Nonpragmatic Vindication of Probabilism.” *Philosophy of Science* 65, 4 (1998): 575–603.

Konek, Jason. “Epistemic Conservativity and Imprecise Credence.” *Philosophy and Phenomenological Research* .

Kraft, Charles H., John W. Pratt, and A. Seidenberg. “Intuitive Probability on Finite Sets.” *The Annals of Mathematical Statistics* 30, 2 (1959): 408–419.

Landes, Jürgen, and Jon Williamson. “Objective Bayesianism and the maximum entropy principle.” *Entropy* 15, 9 (2013): 3528–3591.

Lewis, David K. “A Subjectivist’s Guide to Objective Chance.” In *Studies in Inductive Logic and Probability*, edited by Richard C. Jeffrey. University of California Press, 1980, 83–132.

Luce, R.D., and H. Raiffa. *Games and Decisions: Introduction and Critical Survey.* Dover books on advanced mathematics. Dover Publications, 1957.

McGrath, Matthew. “Being Neutral: Suspension of Judgement, Agnosticism and Inquiry.” *Noûs* .

Milnor, J. “Games Against Nature.” In *Decision Processes*, edited by CH Coombs RM Thrall, and RL Davis. Wiley, New York, 1954.

Myrvold, Wayne C. “Probabilities in Statistical Mechanics.” In *The Oxford Handbook of Probability and Philosophy*, edited by Christopher Hitchcock, and Alan Hajek. Oxford University Press, 2014.

Norton, John D. “A Material Theory of Induction.” *Philosophy of Science* 70, 4 (2003): 647–670.

———. “Probability Disassembled.” *British Journal for the Philosophy of Science* 58, 2 (2007): 141–171.

———. “Ignorance and Indifference.” *Philosophy of Science* 75, 1 (2008): 45–68.

———. “Cosmic Confusions: Not Supporting Versus Supporting Not.” *Philosophy of Science* 77, 4 (2010): 501–523.

Norton, John D. *The Material Theory of Induction.* University of Calgary Press, 2021.

Norton, John D. “Eternal Inflation: When Probabilities Fail.” *Synthese, special edition “Reasoning in Physics,”* .

Peterson, M. *An Introduction to Decision Theory.* Cambridge Introductions to Philosophy. Cambridge University Press, 2009.

Pettigrew, Richard. *Accuracy and the Laws of Credence.* Oxford University Press UK, 2016a.

———. “Accuracy, Risk, and the Principle of Indifference.” *Philosophy and Phenomenological Research* 92, 1 (2016b): 35–59.

Radzvilas, Mantas, William Peden, and Francesco De Pretis. "The Ambiguity Dilemma for Imprecise Bayesians." *The British Journal for the Philosophy of Science* .

Rosa, Luis. "Rational Requirements for Suspended Judgment." *Philosophical Studies* 1–22.

Rott, Hans. *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning.* Oxford University Press, 2001.

Scott, Dana. "Measurement structures and linear inequalities." *Journal of Mathematical Psychology* 1, 2 (1964): 233 – 247.

Shafer, Glenn. *A Mathematical Theory of Evidence.* Princeton University Press, 1976.

Staffel, Julia. "Credences and Suspended Judgments as Transitional Attitudes." *Philosophical Issues* 29, 1 (2019): 281–294.

Tang, Weng Hong. "Reliabilism and the Suspension of Belief." *Australasian Journal of Philosophy* 94, 2 (2015): 362–377.

Titelbaum, Michael. *Fundamentals of Bayesian Epistemology.* Oxford University Press, 2022.

Urbach, Peter, and Colin Howson. *Scientific Reasoning: The Bayesian Approach.* Open Court, 1993.

White, Roger. "Epistemic Permissiveness." *Philosophical Perspectives* 19, 1 (2005): 445–459.

White, Roger. "Evidential Symmetry and Mushy Credence." In *Oxford Studies in Epistemology*, edited by T. Szabo Gendler and J. Hawthorne. Oxford University Press, 2009, 161–186.

Williamson, Jon. "Motivating objective Bayesianism: from empirical constraints to objective probabilities.", 2007.

———. *In Defence of Objective Bayesianism.* OUP Oxford, 2010.

———. "Justifying the principle of indifference." *European Journal for Philosophy of Science* 8 (2018): 1–28.