London School of Economics and Political Science

Beliefs About the Unobserved

Chloé de Canson

A thesis submitted to the Department of Philosophy, Logic and Scientific Method
of the London School of Economics for the degree of Doctor of Philosophy.
London, June 2020

Declaration.

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work.

I confirm that the fourth chapter, "Objectivity and the Method of Arbitrary Functions" is to be published in the *British Journal for the Philosophy of Science* (Canson, forthcoming).

I declare that my thesis consists of 50622 words.

Abstract.

What should one believe about the unobserved? My thesis is a collection of four papers, each of which addresses this question. In the first paper, "Why Subjectivism?", I consider the standing of a position called radical subjective Bayesianism, or subjectivism. The view is composed of two claims—that agents ought to be logically omniscient, and that there is no further norm of rationality—both of which are subject to seemingly conclusive objections. In this paper, I seek, if not to rehabilitate subjectivism, at least to show its critic what is attractive about the position. I show that the critics of subjectivism assume a particular view about justification, which I call the telic view, and that there exist an alternative view, the poric view, on which subjectivism is appealing. I conclude by noting that the tension between telic and poric conceptions of justification might not be an easy one to resolve. In the second paper, "Bayesianism and the Problem of Induction", I examine and reject the two existing Bayesian takes on Hume's problem of induction, and propose my own in their stead. In the third paper, "The Nature of Awareness Growth", I consider the question of how to model an agent who comes to entertain a new proposition about the unobserved. I argue that, contrary to what is typically thought, awareness growth occurs by refinement of the algebra, both on the poric and the telic pictures of Bayesianism. Finally, in the fourth paper, "Objectivity and the Method of Arbitrary Functions", I consider whether, as is widely believed, a mathematical theorem known as the method of arbitrary functions can establish that it is in virtue of systems' dynamics that (some) scientific probabilities are objective. I differentiate between three ways in which authors have claimed that dynamics objectivise probabilities (they putatively render them: ontically interpreted, objectively evaluable, and high-level robust); and I argue that the method of arbitrary functions can establish no such claims, thus dampening the hope that constraints in what to believe about the unobserved can emerge from dynamical facts in the world.

Acknowledgements.

Firstly, I would like to express my deepest gratitude to Anna Mahtani. Her careful reading and probing discussion of every argument in this thesis has allowed me to avoid many falsehoods, and to find my way to the core idea I was trying to express. By believing in me, she has helped me find faith in moments of discouragement. And maybe most importantly, she has served as a model for the kind of philosopher I am trying to become: curious, imaginative, rigorous, and immensely generous. I hope that I have succeeded in acquiring some of these virtues, at least to some extent, and that they transpire in this thesis. To the extent that they do, I have Anna to thank.

Secondly, I would like to thank Richard Bradley and Miklós Rédei, who have also played crucial roles in supervising this thesis. I am incredibly grateful for their time and commitment to teaching me, for the wealth of knowledge that I have acquired from them, and for their scepticism of most of my views. Their reluctance to accept most of the things I have thought in the past few years has forced me to refine and strengthen my arguments, which has led to a deep improvement of my work.

Finally, although Arif Ahmed has not read a word of this thesis, he must be credited for infecting me with inductive scepticism. His refutation of a passing remark on induction in a footnote of my MPhil thesis, in the Faculty of Philosophy at the University of Cambridge, has served as the trigger for my interest in the topic. Although I was used to Arif's refutation of all the ideas I proudly came to present, the refutation of this specific footnote allowed me to appreciate the problem of induction in its full daunting strength, and it is now clear to me that this conversation with him has been the sprout of this thesis.

So many people have read ancestors of the papers in this thesis, attended talks at which I presented some of the ideas contained in them, and provided feedback on them. For their invaluable contribution to parts of this thesis, I would like to thank Marshall Abrams, Liam Kofi Bright, Jeremy Butterfield, Joshua Eisenthal, Leah Henderson, David Kinney, Bryan Roberts, Jan-Willem Romeijn, Joe Roussos,

Michael Strevens, Mauricio Suárez, and Justin Vlasits. I would also like to thank three anonymous referees for the *British Journal for the Philosophy of Science*, whose comments on the fourth chapter have prompted considerable improvements.

The Department of Philosophy, Logic, and Scientific Method at LSE has been an incredible place to work. Thanks to my fellow PhD students for their friendship and intellectual generosity, in particular to Dmitry Ananyev, Nicolas Côté, Christina Easton, Margherita Harris, Sophie Kikkert, David Kinney, Ko-Hung Kuan, Nick Makins, Silvia Milano, Joe Roussos, Bastian Steuwer, and Cecily Whiteley. The PhD community has been a great source of learning, support, and joy, and I will miss it dearly.

Thanks also to many members of the faculty who, despite being officially unrelated to me and my thesis, have provided philosophical, practical, and personal support to me in the past four years. Thanks in particular to Jason Alexander, Jonathan Birch, Liam Kofi Bright, Laurenz Hudetz, Christian List, Bryan Roberts, and Johanna Thoma.

Thanks to my students over the years, for teaching me so very much, in particular to Carlo Bellomo, Jan-Paul Sandmann, and Val Weisner.

Thanks finally to the professional services staff members, and in particular to Andrea Pawley for her help in this past year.

Thanks to the workers of the British Library for a perfect writing environment.

Thanks to the poets and writers whose words carried me through.

Finally, thanks to those who have shaped me through their love. To my family: Sophie, Olivier, Agathe, and Adèle; and to Aline, Anna, Anne, Ariane, Auguste, Caroline, Juliette, Laure, Nicolas, and Pauline. In loving memory of two of us who died while I was writing this thesis: Raoul, and Ka—who, on her deathbed at the age of 103, called philosophy "a science for the mad".

To my friends outside of academic philosophy, in particular to Ana, Chris, Diane, Fabian, Maïa, Oli, and Thomas.

Lastly and most of all, for their love, to Camilla, Dimi, Fee, Indy, and Matthew.

Contents.

Figures and Tables.

...the fact that we all just carry on as if everything's going to be okay, but we don't *know* that, the fact that you're just bumbling along and at any moment a flower pot could fall on your head and leave you in a wheelchair from then on, or you could just go blind for no reason, or the doctor could tell you that you have a year to live, [...] the fact that so many things can go wrong, and *so* wrong, the fact that you never know what's coming down the pike, [...] the fact that it's amazing we're not *crippled with dread* all the time, about what the future will bring, [...] the fact that I cling to the few certainties there are about it, the fact that some things are definite, like

the sun will rise and set every day
without fail
absolutely without fail
the moon will come and go too
though fewer and fewer people
including me
will understand its phases
bread will rise
children will grow
eBay will flourish
Facebook likewise...

Lucy Ellmann, *Ducks, Newburyport*

0. Introduction

A Question

Every time I have eaten bread, it has nourished me. Every time I have boarded a plane, it has flown safely. Every time I have taken some paracetamol, it has calmed the pain. Of course, I cannot be absolutely certain that these things will continue—bread to nourish, planes to fly, painkillers to calm pain. After all, my evidence does not rule out the possibility that they might stop. But, the thought goes, although my evidence does not rule out this possibility, it does make it unlikely; and so I ought to be very confident that they will continue. Is this thought correct? Is it indeed the case that I am justified in my high degree of confidence that bread will continue to nourish, that planes will continue to fly, and that paracetamol will continue to soothe?

The Plan

To answer this question, one must ask: what might we mean when we say that a particular credence is *justified*? A popular answer comes from Frank P. Ramsey, probably the single greatest influence on Bayesian thought. In his seminal paper, "Truth and Probability" (1926), he argues that agents' degrees of belief ought to satisfy particular norms, by showing that if they do not, things will turn out badly for them in some way. Thus Ramsey binds the rationality of particular epistemic attitudes with the consequences of having these attitudes: the rightness of epistemic states is determined by the goodness of their consequences. This *telic* approach to justification has become pervasive in Bayesian epistemology.

In the first chapter of this thesis, I propose an alternative. In a spirit akin to Hume's (1739, 1748), I suggest that we attend to the means by which agents conduct their inquiry. This intimates a different account of epistemic justification, on which a particular epistemic state is rational just in case it is warranted by the agent's means of inquiry. I call such an approach *poric*, after the Greek πόρος, which refers to the means to one's ends.[1] I then show that, on this approach, no credence towards a proposition about the unobserved is more justified than any other. So, I answer *no* to the central question: we can never be justified in having a high (rather than low) credence that bread will nourish, planes will fly, and painkillers will soothe.

The remaining three chapters can be read as elaborations on different aspects of this first chapter. In the second chapter, I examine the relationship between Hume's problem of induction and Bayesianism in general. In the third chapter, I examine the consequences of an approach to epistemology I carve out in the first chapter, for the question of how Bayesians can and ought to model awareness growth. And in the fourth chapter, I examine the philosophical relevance of the method of arbitrary functions, picking up a thread in the first chapter on how agents might come to know the values of objective chances. As the reader will notice, the first three chapters are interconnected papers, which build on one another but can be read in isolation; while the fourth is more separate. In the rest of this introduction, I provide an extended abstract of each.

1. Why Subjectivism?

Few positions are as unpopular as radical subjective Bayesianism, the view that no (coherent) credences about the unobserved are any more justified than any (coherent) others. Indeed, as far as I know, this view has not been defended in print by any philosopher: even the classic supposed subjectivists—Ramsey, de Finetti, and Jeffrey—shy away from the view. In this paper, I seek, if not to rehabilitate

[1] The word also refers to a passageway, especially a passage over a body of water, such as a bridge; and derivatively, to a journey or crossing.

subjectivism, at least to show its critic what is attractive about the position.

I argue that what is at stake in the subjectivism/anti-subjectivism debate is not, as is commonly thought, which norms of rationality are true, but rather, the account of rationality that we adopt. Indeed, although there are good and well-rehearsed arguments against subjectivism on the widespread telic approach to rationality (accounting for the view's unpopularity), things are not so on the poric approach. I put forward a positive poric argument for subjectivism, and I show that the typical complaints against the view—that, to the extent that it mandates particular forms of omniscience, it is too demanding; and to the extent that this is all that it mandates, it is not demanding enough—cannot be articulated plausibly on the poric approach. It follows that the debate between subjectivists and anti-subjectivists is best understood as a debate between a poric and a telic account of rationality.

I conclude the paper by examining the prospects for the poric/telic debate. I submit that, if we accept a conception of epistemology on which agents inquire with the intention of determining what is the case—what I call the *means-ends* conception—the poric account is a strong contender as an approach to justification. Indeed, it directly binds the agent's means of inquiry with the epistemic state it would be rational for her to have. This is bad news for telic Bayesians: they must either argue that telism is a better approach to means-ends rationality, if they want to retain the means-ends conception of epistemology; or reject the means-ends conception of epistemology altogether, in which case they must provide an alternative account of what they are doing, and why it matters.

2. Bayesianism and the Problem of Induction

In this paper, I examine the relationship between Bayesian epistemology and Hume's problem of induction. I identify two existing strands of thought on this relationship. On the one hand, the *inferentialists* take Hume's problem to be that of supplying a sound inductive logic—the logic of inductive inference. They argue that, since Bayesianism provides exactly that, it constitutes a solution to Hume's

problem. On the other hand, the *superbaby theorists* take Hume's problem to be that of justifying an agent's credences at the time of analysis. They argue that this would involve providing an account of which credences an agent ought to have had prior to having any evidence, and an account of how she ought to have updated these credences as she subsequently received evidence; and since Bayesianism has supplied the latter, it constitutes partial progress towards a solution to Hume's problem, but insofar as it has not supplied the former, this solution is incomplete. Thus the two camps disagree on what the problem is, and on whether Bayesianism constitutes a complete solution to it, but they both agree that, through their defence of conditionalisation, Bayesians have made progress on the problem.

I agree with the superbaby theorists, against the inferentialists, that the problem of induction is that of justifying agents' epistemic states: the problem is epistemic, not logical, and which is the correct logic of inductive inference (if there is such a thing) matters only insofar as it helps with the question of what to believe. But I argue that the superbaby theorists are misguided in two respects. I show that, given conditionalisation, justifying the relevant credences at any time justifies them at every other time, and so there is nothing special about superbaby credences. But more importantly, I argue that justifying credences about the unobserved is hard for the same reason regardless of time, including when the agent is a superbaby and at the time of analysis. In light of this, the problem of induction is more general than superbaby theorists take it to be: it is that of justifying credences about the unobserved.

Four things of interest follow from my arguments. The first is that, contrary to what both inferentialists and superbaby theorists contend, conditionalisation does not constitute progress toward a solution to Hume's problem. The second is that Bayesians do not need to appeal to superbabies, and so avoid having to argue that this particular kind of idealisation is justified. The third is that a widespread strategy to avoid engaging with the problem of induction—namely, insisting that all depends on the problem of the priors, as if it was a separate problem—is infelicitous; the problem cannot be pushed back thusly. The fourth is that my analysis

positions subjectivism, objectivism, and other putative alternatives as direct responses to the problem of induction. This entails, I argue, that subjectivism asserts inductive scepticism.

3. The Nature of Awareness Growth

Real agents often undergo what is known as *awareness growth*: they come to consider propositions of which they were previously unaware. Not only does such an epistemic shift not seem irrational, but it appears to be a crucial part of epistemic thriving. As such, it is important for Bayesians to accommodate this phenomenon. Two proposals for how to do this have been proposed: I call them the *expansion* view and the *refinement* view. In the developing literature on this topic, there are no outright defendants of the refinement view: not only are there critics, but even those who endorse it do so half-heartedly. This is because of two universally held claims. The first, I call the *contrivance* claim: the refinement view, and only the refinement view, mandates the inclusion of a catch-all proposition in the algebras of less-than-fully-aware agents. The second, I call the *defectiveness* claim: catch-all propositions are defective, in the sense that it is uniquely difficult to assign a justified credence to them. Taken together, these claims constitute an argument against the refinement view.

In this paper, I argue resolutely in favour of the refinement view. Firstly, I argue against the contrivance claim. I begin by considering a widely discussed putative norm of rationality, *regularity*. I consider two arguments against it, and although I claim that one of them can be set aside, the other should move us to adopt a weaker claim instead, which I call *humility*, according to which agents ought not have extremal credence in propositions about the unobserved. Humility entails that the algebras of less-than-fully-aware agents ought to include a catch-all proposition. So, I conclude, the contrivance claim is false: regardless of whether agents ever undergo awareness growth, let alone whether they do so by refinement, the inclusion of a catch-all proposition is mandated. I finish this part of the paper by considering a way some have sought to reject the inclusion of a catch-all

proposition without rejecting the spirit of humility, and I show that, in order to make sense of this kind of position, one has to adopt a very unattractive account of credence.

Secondly, I argue against the defectiveness claim. I reconstruct the argument for this claim as consisting of two steps: (1) the content of a catch-all proposition is opaque, and (2) this makes it uniquely difficult to assign a credence to. I rebut this argument in three steps: I shed doubt on the claim that the transparent/opaque distinction can be meaningfully drawn; I show that even if it can, it cannot be the distinguishing mark of the catch-all; and I show that whatever makes the catch-all supposedly special cannot make it uniquely difficult to justify credences in. I then propose a disjunctive, positive argument against the defectiveness claim: I show that whether one adopts a poric or a telic account of rationality, one ought to reject it.

It follows that the reasons to reject the refinement view are a lot less convincing than they are usually taken to be. I then consider the expansion view and the refinement view side by side, and I conclude in favour of the latter.

4. Objectivity and the Method of Arbitrary Functions

The method of arbitrary functions has attracted a lot of attention, because it is widely believed to provide a reason for thinking that (at least some of) the probabilities in the special sciences are objective in virtue of the dynamics of the systems under study. More precisely, many take it to show that the dynamics of a system can serve to objectivise an underlying probability function; or in other words, that the objectivity of the probabilities can emerge from the dynamics of the system. In this paper, I argue that this is not the case. More precisely, I argue that, to the extent that the aforementioned probabilities are objective, they are not so in virtue of dynamics *qua* dynamics. Thus my argument is consistent with the claim that the probabilities in question are objective; and it is even consistent with the claim that the method of arbitrary functions can shed some light on the question of why they are objective (though for this latter claim, not in all senses of *objective*).

But, I show, this has nothing in particular to do with dynamics, rather at best, dynamics are an instance of a much more general way in which the probabilities in question could be objectivised.

I begin by differentiating three ways in which a probability function might be thought to be objective. A probability function might be ontically (as opposed to epistemically) interpreted, if it represents a mind-independent phenomenon. If it is interpreted as a credence function, it might be objectively (as opposed to subjectively) evaluable, if disagreement about its values entail fault. Finally, it might be high-level robust (as opposed to chaotic) if the values it assigns to high-level phenomena do not depend much on the values it assigns to low-level phenomena. I then show that, for each of these senses of *objective*, some of the authors in the literature have argued that the method can be used to show that the dynamics of a system can establish the objectivity of probability functions in the sense in question. I argue that this is a mistake: the method does not show that the objectivity of probability functions stems from the system's dynamics, in any of the senses of objectivity that can be found in the relevant literature.

I conclude that the method of arbitrary functions does not fulfil the hope of locating the objectivity of scientific probabilities in the mechanics of the phenomena associated with these probabilities. But, the method does have philosophical value nonetheless: in particular, it can be used to determine when one ought to be careful about replicating initial conditions precisely in probabilistic experiments.

*

As the reader will have noticed, this last chapter is less tightly connected to the first three than they are amongst themselves. What motivated me to think about the method of arbitrary functions was the hope that facts about objective probabilities could emerge from purely empirical phenomena—such as the mechanics of wheels of fortune. This, as I discuss in "Why Subjectivism?", would have explained how agents with the capacities to observe and reason, but no more, could come to know chances. This in turn, would have provided more stringent constraints on

rationality than subjectivism allows. Unfortunately, as I argue in this paper, I do not think that the method shows such a thing.

1. Why Subjectivism?

Abstract. Few philosophical positions are as unpopular as radical subjective Bayesianism. In this paper, I seek, if not to rehabilitate subjectivism, at least to show its critic what is attractive about the position. I argue that what is at stake in the subjectivism/anti-subjectivism debate is not, as is commonly thought, which norms of rationality are true, but rather, the conception of rationality that we adopt: there is an alternative approach to the widespread telic approach to rationality, which I call the poric approach, on which subjectivism is an attractive position. I end by considering the prospects for escaping subjectivism, and I suggest that they are bleak.

0.

One would be hard-pressed to find a self-proclaimed radical subjective Bayesian (henceforth, *subjectivist*)[1] today, and with good reason: the view suffers from serious problems. Critics of the position point out that, to the extent that it mandates particular forms of omniscience, it is too demanding; and to the extent that this is all that it mandates, it is not demanding enough. In light of these problems, subjectivism has been widely rejected.

[1] The term *subjectivism* unfortunately has two meanings in the philosophy of probability. It refers both to the epistemic interpretation of probability functions, as opposed to the various ontic interpretations; and to a particular view about which epistemic probabilities are justified. In this paper, I am concerned with the latter.

In this paper, I seek, if not to rehabilitate subjectivism, at least to show its critic what is attractive about the position. But my strategy won't be to engage head-on with the anti-subjectivist. Rather, I will highlight that many anti-subjectivists adopt a particular account of epistemic justification—one which might in principle be rejected. I will then propose an alternative account that one might adopt in its stead, and show not only that the serious problems with subjectivism lose their traction on this new account; but also that a strong argument for subjectivism emerges out of it. Thus I will carve an alternative view of rationality on which—I will argue—subjectivism holds. Although I will not dispute the standing of anti-subjectivism on the original account of rationality, I will end this paper by suggesting that the account I propose is better suited to play the role that the dominant account is supposed to play. This tells in favour of my account, and thus, of subjectivism.

The plan for the paper is as follows. I begin with a presentation of subjectivism and its two usual objections (§1). Then, I articulate the account of justification assumed across Bayesian epistemology, and I propose an alternative to it (§2). I construct an argument for subjectivism, based on this alternative account (§3), and I show that the two usual objections fail if it is accepted (§4–5). I conclude by examining the prospects for escaping subjectivism, and I suggest that they are bleak (§6).

1.

The consideration at the heart of Bayesian epistemology is that beliefs come in degrees: agents may be more or less confident in the truth of any particular proposition.[2] The agent's degree of belief in a proposition is called her *credence* in that proposition, and her total epistemic state can be represented by a credence function $p : \mathcal{A} \to [0, 1]$, which assigns a credence to each element in the set $\mathcal{A}$ of all propositions the agent entertains. The formal structure of $\mathcal{A}$ is as follows. We begin with a non-empty set Ω. We then construct $\mathcal{A}$ to be a Boolean algebra of

[2] I use the term *Bayesian* in a broad way to pick out any theory or theorist concerned with agents' credences.

Ω, which is to say that it is a set of subsets of Ω, such that $\Omega, \emptyset \in \mathcal{A}$, and $\mathcal{A}$ is closed under union and negation. I leave the question of how to precisely interpret these mathematical objects for later on. However I want to draw the attention of the reader to the distinction between the propositions we might call *trivial* (Ω and $\emptyset$), and those we might call *non-trivial* ($A_1, A_2, ...$). This distinction will turn out to be important.

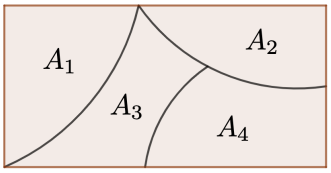

1. An algebra of propositions

Much of Bayesian philosophical scholarship is organised around the question of which norms, if any, govern an agent's credences on such sets of propositions. Subjectivism is a particular view on the matter:

Subjectivism. There are two (and only two) credal norms of rationality:

Probabilism. An agent's credence function p must be a probability function, that is, such that:

(a) $p(\Omega) = 1$ and $p(\emptyset) = 0$.

(b) $p(A_i \vee A_j) = p(A_i) + p(A_j)$ for all inconsistent propositions A_i, A_j.

Conditionalisation. Upon receiving some evidence E, an agent's new credence function p' must be such that:

(a) $p'(E) = 1$ and $p'(\neg E) = 0$.

(b) $p'(A_i|E) = p(A_i|E)$ for all propositions A_i.

This somewhat unorthodox presentation of subjectivism highlights the fact that each of the norms that compose it are made of one sub-norm **(a)** we might call *substantive*, which prescribes particular credences to particular propositions; and of another sub-norm **(b)** we might call *formal*, which makes particular coherence

prescriptions on the entire credence function(s).[3] These can be presented in the table below.

	Probabilism	**Conditionalisation**
Substantive	Trivial Omniscience	Evidential Omniscience
Formal	Additivity	Rigidity

2. Sub-norms of subjectivism

In this paper, I shall be concerned solely with the substantive sub-norms of subjectivism, for two reasons. The first is that, as we shall see, they are the ones at stake in the dispute between subjectivists and their opponents. The second is that formal sub-norms require an altogether different treatment, different enough that it can be neatly treated separately. As a result, I leave the question of formal sub-norms for future work.

In light of this restriction, we can consider a restatement of subjectivism.

Subjectivism. An agent is (substantively) rational if and only if she is

Trivially omniscient. $p(\Omega) = 1$ and $p(\emptyset) = 0$.

Evidentially omniscient. $p(E_i) = 1$ and $p(\neg E_i) = 0$, for all evidential propositions E_i.

And this restricted claim may be restated in a different form once again:

Subjectivism.

Necessity claim. Agents ought to be trivially and evidentially omniscient.

Sufficiency claim. There are no other (substantive) requirements of rationality.

With such a statement in hand, we are in a position to see why subjectivism is so widely rejected: there are what seem to be decisive arguments against both the necessity and the sufficiency claims. But before reviewing these arguments, let

[3] I call these sub-norms *substantive* and *formal* because I think that these labels capture nicely the way in which they differ, but nothing of substance hangs on this. If the reader rejects this characterisation of the distinction, she will (I hope) still accept the restriction of my concern to one of the disjuncts.

us first clarify what is meant by trivial and evidential omniscience. Let us begin with trivial omniscience. Without an interpretation of the agent's algebra, we do not know what Ω and $\emptyset$ represent, and so we cannot know what it means to be omniscient with respect to them. Now, the problem with the necessity claim in the literature is often called the *problem of logical omniscience*,[4] and correspondingly, probabilism is typically taken to entail that agents ought to have credence 1 in all logical truths and credence 0 in all logical falsehoods. This shows that Bayesians typically interpret the formalism in such a way that logical truths and logical falsehoods—tautologies and contradictions—should be modelled by the trivial elements. So, we shall assume that trivial omniscience (at least) entails logical omniscience. (More on this assumption later.)

What about evidential omniscience? Again, without an interpretation of the algebra, we are not in a position to determine what qualifies as an evidential proposition. This is because, in order to determine what evidential omniscience requires of agents, we must adopt a specific conception of evidence. Now, there are two options. On the *externalist* conception of evidence, an agent's evidential propositions express facts about the external world, facts such as those expressed by the sentences "the sun has risen today" and "it is a book I am holding". On the *internalist* conception of evidence by contrast, evidential propositions express facts about the agent's internal world, facts such as those expressed by sentences like "it seems to me that the sun has risen today" and "it seems to me like it is a book that I am holding". For now, it will suffice to remark that most Bayesians (explicitly or implicitly) assume an externalist conception of evidence. Typical examples of evidence in the literature include things like: the coin has landed heads, and there

[4] In fact, what is usually called "the problem of logical omniscience" actually refers indiscriminately to the problems associated with two distinct requirements: (a) that agents be certain of the truth-value of all logical propositions; but also (b) that agents' credences in different propositions be related in particular ways, for instance that an agent always be more confident in a proposition X's logical consequence than in X. This amounts to treating substantive and formal components of probabilism at once. But as I announced earlier in the paper, I will restrict all my claims to substantive issues; in this case, to (a).

is a goat behind the door.

Let us begin with the necessity claim: what is the problem with mandating logical and evidential omniscience? In order to have credence 1 or 0 in all logical and evidential propositions, the argument goes, one must be in a position to recognise them as such. But real human agents are not in a position to do this. Indeed, not only are there infinitely many logical truths, some of which very complex, but many have been shown not to be determinable even in principle. So many human agents, when faced with a logical proposition, are not in a position to tell that they are. Furthermore, real human agents are not always in a position to recognise what evidence (externalistically conceived) they possess. For instance, an agent may not be able to tell whether the proposition expressing that she has hands is an evidential proposition for her. So many human agents, when faced with an evidential proposition, are not in a position to tell that they are. It follows that agents cannot be omniscient in the ways mandated by subjectivism: this is the over-demandingness objection.

Over-demandingness objection.

1. Agents ought to be trivially and evidentially omniscient. (Necessity Claim)

2. Ought implies can. (Premise)

3. Agents can be trivially and evidentially omniscient. (1, 2)

4. They cannot. (Premise)

5. So, the necessity claim is false.

In light of this objection, most have rejected the necessity claim as interpreted above. Now, because of a widespread reticence to abandoning probabilism (the reason for which I discuss in fn. 11), the main way in which Bayesian scholars have done this is not by rejecting trivial omniscience, but by reinterpreting it in such a way as to undermine premise **4**. For example, Hacking (1967) and more recently Pettigrew (forthcominga), and in a different way Bradley (2017a), have proposed reinterpretations of the trivial/non-trivial distinction on which trivial omniscience no longer entails logical omniscience. In that way, the fact that real human agents

cannot recognise every logical truth no longer challenges the mandate of trivial omniscience.

An analogous strategy has been attempted to ward off the challenge from evidential omniscience: some have suggested reinterpreting the evidential/non-evidential distinction in an internalist way, so that evidential propositions express *seem*-facts and agents are always in a position to be certain of their evidence.[5] This strategy has not been very popular however. As Jeffrey (1968) puts it: "for excellent reasons, this move is now in low repute" (p. 171). He argues that we should reject evidential omniscience, and replace it with a norm according to which, when agents acquire evidence, they do not adopt credence 1 in E and credence 0 in $\neg E$, but instead come to have (probabilistic) non-trivial credences in these propositions. In that way, agents can come to assign credence .9 to the proposition that they are holding a book, and credence .1 to the proposition that they are not (because, perhaps, they are being deceived by an evil daemon).[6] More recently, rejections of

[5] An in-depth examination of the reinterpretation strategies is beyond the scope of this paper. However, it will be relevant to note that such reinterpretations stand in tension with what is often taken to be the target-subject of Bayesian epistemology; namely, inquiry into the empirical world. By virtue of the formalism, the propositions in which we are interested—in which we want to know what credence to have—are the non-trivial propositions. Furthermore, by virtue of the formalism too, each of these propositions is liable to become evidence. In light of this, the Bayesian enterprise is naturally understood in such a way that non-trivial propositions are *empirical* propositions, that is, propositions about the external world that might come to be known by observation (capaciously understood). But reinterpretations of the algebra which allow non-trivial propositions to be logical propositions or propositions about the agent's mind stand in tension with this. For on these reinterpretations, logical propositions could in principle constitute evidence, and propositions about the external world could not. Clearly, my claims here do not constitute a definitive refutation of reinterpretation strategies. They do however call for an answer to the question of which domain(s) of inquiry the proponents of these startegies take Bayesian epistemology to be concerned with.

[6] Jeffrey retains rigidity (the formal sub-norm of conditionalisation), and ends up endorsing a new updating norm called *Jeffrey conditionalisation.*

conditionalisation motivated by the failure of evidential omniscience can be found in Bronfman (2014), Schoenfield (2017a), and Gallow (2017, forthcoming). To sum up: the over-demandingness objection is widely taken to be significant, and to call for a weakening of the necessity claim.

Let us now move on to the objection to the second sub-norm of subjectivism: the sufficiency claim. This claim has been the target of the vast majority of anti-subjectivist criticism, to the point that the subjectivist/anti-subjectivist debate is often presented in those terms only. What the sufficiency claim asserts is that there is no substantive requirement of rationality beyond trivial and evidential omniscience. In other words, as long as agents are omniscient in those ways, they are rational. So, take an arbitrary non-trivial non-evidential proposition: any credence in this proposition is rational. But surely that cannot be, the argument goes. Surely someone who has an extremely high credence that a coin will land heads, despite knowing nothing about the coin in question, is irrational. So, subjectivism is not sufficiently restrictive: this is the under-demandingness objection. Calling a coin about which one has no evidence an *unknown coin*, the objection goes as follows:

Under-demandingness objection.

1. All that's rationally required of an agent is for her to be certain in the truth of trivial and evidential propositions. (Sufficiency Claim)

2. A credence of .99 in the proposition that an unknown coin will land heads is rational. (1)

3. It is not. (Premise)

4. So, the sufficiency claim is false.

In response to this objection, many have rejected the sufficiency claim, and proposed variously strong norms to supplement probabilism and conditionalisation. In one tradition, people have attempted to constrain credences by appealing to norms according to which agents' credences are required to match exogenously given probabilities. In another (sometimes overlapping) tradition, people have proposed indifference norms on credences, which prescribe a particular unique

credence to each proposition in the algebra in the absence of evidence. In yet another (sometimes also overlapping) tradition, people have appealed to evidentialist considerations by insisting that credences must reflect the evidence. I will discuss these traditions in more detail in §5.

This objection has turned out to be particularly potent in deterring people from subjectivism. Indeed, not only are there, to my knowledge, no contemporary defences of the sufficiency claim in print, but there are also reasons to believe that even the original subjectivists—Ramsey, de Finetti, and Jeffrey—rejected it. Ramsey says: "if we are told that one of these people's names begins with A and that there are 8 of them, it is reasonable to believe to degree $1/8^{\text{th}}$ that any particular one's name begins with A, and this is what we should all do (unless we felt that there was something else relevant)" (1926, pp. 100–101). Jeffrey writes that, as far as the assignment of credences is concerned, his view is "often faulted as uncritical [...]; 'anything goes'. But the adoption of [credences] is a subject-matter dependent *techne*, an art of judgment [...]. Although [the expert agent] is far from knowing how [her mechanism for assigning credences] works, she can know *that* it works, pretty well" (1991, pp. 11-12).[7] And even de Finetti says: "though maintaining the subjectivist idea that no fact can prove or disprove belief, I find no difficulty in admitting that any form of comparison between probability evaluations (of myself, of other people) and actual events may be an element influencing my further judgment, of the same status as any other kind of information." (1962, p. 360).[8]

Given the strength of both the over- and under-demandingness objections, and the attractiveness of many available alternatives, it should come as no surprise that subjectivism enjoys a low level of popularity. Thus, although subjectivism tends to be viewed as the default position in Bayesianism, this is not because it is taken to be true. Rather, I submit, it is because it is taken to be simple; and so it is discussed, not in the sense of being engaged with, but in the sense of being used

[7] So it turns out that Jeffrey rejects both the necessity and the sufficiency claim!

[8] For a history and critical discussion of the early subjectivists' commitments, see Galavotti (2011, 2016, 2018).

to define putatively more plausible positions by contrast. My aim in this paper is to counter this trend, and to provide an argument *for* subjectivism. I will present a way of thinking about justified credence on which subjectivism is an appealing answer to the question of what norms govern epistemic attitudes. I will then show that the over-demandingness and the under-demandingness objections fail to have any grip on this alternative way of thinking about justification. Finally, I will take seriously the thought that it is in some sense bad if subjectivism is true, but will suggest that, in light of my argument, the prospects for escape are bleak.

2.

As I remarked in the previous section, Bayesians have organised their field of study around the question of which graded epistemic attitudes are justified. What I want to show in this section is that many Bayesians have been taking a particular interpretive approach to this question; and that there exists an other one might take. For now, I shall call them the *standard approach* and the *alternative approach* respectively.

Let me begin by sketching the standard approach. Bayesians working within this approach are ultimately interested in questions like: What is the best credal state to have for a climate scientist in this (specified) situation? What is the best credal state to have for a person who wonders whether it is about to rain? These questions are interesting to the standard Bayesian because of the two characteristics he takes epistemic states to have. On the one hand, credences aim to represent the world; on the other, they guide rational action. So, the standard Bayesian wants to know whether the scientist should have a high credence that temperatures will rise because he wants to know whether she is correct in her confidence that they will, and because he wants to know what she (and therefore we) should do in response to the climate crisis. So, determining what it is best to believe, and thereby (given fixed desires) what it is best to do, is at the heart of the standard Bayesian project. Given this broad purpose, his methodology tends to look as follows. He starts with a (real or imagined) human agent, such as the climate scientist or the rain-

averse pedestrian. He then isolates the normatively relevant aspects of the agent's epistemic state: her credences; which he represents, or *models*, using a credence function.[9] Then, he asks: What shape should this credence function have? Or equivalently, which credences are justified? By this, he means: Which credences would it be best for the agent to have, given the aim of credences?

Here is an alternative approach to Bayesian epistemology. On this approach, the question in which the theorist is ultimately interested is: which justified epistemic states can our means of inquiry afford us? In other words: What justified credences can we achieve on the basis of our means of inquiry? How successful can we be at finding out about the world? In order to address these questions, the alternative Bayesian adopts a different methodology than the one described above. She begins by listing the agent's means of inquiry. These might be reason, observation, testimony, introspection, and so forth. Then, for each proposition that the agent considers, she asks: do the agent's means of inquiry warrant a particular credence in this proposition? For example, has the agent observed that it is true? Has she been told that it is true? And so forth. More generally: on the basis of her means of inquiry, what kind of epistemic attitudes towards these propositions are warranted?

These two approaches differ in at least one salient way: they differ in their account of justification. On the standard approach, a credal state is justified just in case it is the best credal state for the agent to have; but on the alternative approach, a credal state is justified just in case it is warranted by the agent's means of inquiry. Thus we can call the standard approach *telic*, after the Greek τέλος, which refers to an end or aim: on it, to be justified is to be related in a particular way to the end, aim, or goal of epistemic inquiry. By contrast, we can call the alternative approach *poric*, after the Greek πόρος, which refers to the means to one's ends: on it, to be justified is to be related in a particular way (namely, warrant) to the means of inquiry.

The poric account. An epistemic state is justified just in case it is

[9] For an investigation into modelling in Bayesian methodology, see Roussos (ms).

warranted by the agent's means of inquiry.

The telic account. An epistemic state is justified just in case it is related in the right way to the aim of inquiry.

On the poric approach, one might say: you ought to have credence 1 in a particular proposition because you have observed that this proposition is true. On a telic approach, one might say: you ought to have credence 1 in a particular proposition because you would be guaranteed to lose money if you did not. We can see from this example that the two approaches can in principle be extensionally equivalent. Indeed, both the telic and the poric theorists rule that the fictional agent above ought to have credence 1 in the proposition in question, albeit for different reasons.

The next three sections will function, among other things, as an illustration and precisification of poric Bayesianism. But in the rest of this section, I will show that the telic approach is widespread in Bayesian epistemology, and I will comment on how the telic/poric distinction relates to epistemic consequentialism and evidentialism. Before I start, a terminological note is in order. I have defined justification on the poric account as consisting of a particular relationship between epistemic states and the means of inquiry, namely, warrant. By contrast, I have left the relationship between justification and the end of inquiry indefinite in my exposition of the telic account. This is because there are many ways of filling in the account; or in other words, there are many possible telic accounts of justification. In what follows, I will sometimes speak correctly and say of a particular account that it is telic, or of a particular argument that it presupposes a telic account of justification, but I will also sometimes talk of *the* telic account. This will be shorthand for the most common account of rationality in the Bayesian literature, which, as I am about to show, is telic. To insist once again: although there is a single poric account, there are (at least in principle) several telic accounts.

As the reader who has made it this far will know, the most popular types of arguments in Bayesian epistemology are given in the table below,[10] together with

[10] For an overview and critical discussion of: Dutch Book arguments, see Pettigrew (2020); accuracy dominance and expected accuracy arguments, see Pettigrew (2016a) and Petti-

the applications of these argument-types to defences of probabilism and conditionalisation.[11]

	Probabilism	**Conditionalisation**
Dutch Book	Ramsey (1926), de Finetti (1937)	Teller (1973), Lewis (1999)
Accuracy Dominance	Joyce (1998)	Briggs and Pettigrew (2020)
Expected Accuracy	Leitgeb and Pettigrew (2010a,b)	Greaves and Wallace (2006)
Decision Theoretic	Ramsey (1926), Savage (1954)	Savage (1954)

3. Arguments for probabilism and conditionalisation

As I will show, these arguments share a common structure: they conform to the following recipe. **(1)** Take a putative norm of rationality X; for instance, probabilism. **(2)** Show that, if an agent's credences violate X, there is an alternative, better credence function that the agent could have had, which satisfies X. **(3)** Conclude that an agent's credences must conform to X. It follows that the proponents of these arguments adopt a telic conception of justification: in virtue of their shared structure, they link up justification to what it would be best for the agent to believe. Indeed this is made explicit by Pettigrew, who calls his position "teleological" (p. 11, 2016a).

Let us begin with the Dutch Book argument and the two accuracy arguments. Firstly, let us consider the Dutch Book argument.[12] It was first formulated by Ramsey (1926), and goes as follows. **(1)** Consider the norm of probabilism. **(2)** Show that, if an agent fails to conform to probabilism, she will accept a series of bets such that she is guaranteed to lose money, no matter how the bets are resolved. **(3)** Conclude that credences ought to be probabilistic. It is plain that the Dutch

grew (2019); arguments for conditionalisation, see Pettigrew (2020).

[11] Bayesians have been much more convinced by arguments for probabilism than they have been by the corresponding arguments for conditionalisation. As a result, probabilism is more widely endorsed than conditionalisation. This might explain why people have been quicker to reject evidential than trivial omniscience!

[12] The etymology of the term is unknown, and a better name for it might be the *utility dominance argument*.

Book argument follows the above recipe. As we shall see, the accuracy arguments do too. The accuracy dominance argument is very closely related to the Dutch Book argument; it differs in its step **(2)**, as it shows that if an agent fails to conform to probabilism, there will be an alternative probabilistic credence function which is guaranteed to be more accurate, no matter how the world turns out to be. The expected accuracy argument is again closely related; it differs in its step **(2)**, which shows that if an agent fails to conform to probabilism, she will fail to minimise the expected inaccuracy of her credences. These three types of arguments presuppose a telic account of justification: it is because the agent will be (pragmatically or alethically) badly off if she doesn't that her credences must satisfy probabilism. (This presentation of the arguments allows us to give a nice characterisation of the relationship between Dutch Book arguments and accuracy arguments: they disagree about what has final value—Dutch Books: utility; accuracy arguments: accuracy—but they agree on the way in which what it is rational for an agent to do relates to what has final value.)[13]

Let us now turn our attention to decision theoretic arguments. These arguments also originate with Ramsey (1926), who proposed them as an alternative to the Dutch Book argument, which he took to have serious flaws. We shall see that they too follow the three-step recipe laid out above, though in a slightly more complicated way. **(1)** Consider the norm of probabilism. **(2a)** Prove what is known as a *representation theorem*, which shows that an agent's preferences over a set of outcomes satisfy particular constraints Y if and only if that agent can be represented as (among other things) having probabilistic credences.[14] **(2b)** Show that, if the agent does not satisfy constraints Y, things will go badly for her in some sense.

[13] Note that many Bayesians believe that credences have dual aims: that of accurately representing the world, and that of adequately guiding actions. This gives them reason to accept both types of arguments.

[14] There are a number of different representation theorems in the decision theory literature. The classic ones which establish that credences must be probabilisitc are: Ramsey (1926), Savage (1954), and Jeffrey (1965) and Bolker (1966); and some more recent influential ones include Joyce (1999), Buchak (2013), and Bradley (2017a).

For instance, all major representation theorems are such that among constraints Y is the constraint that preferences be *acyclical*; that is, not such that $a \prec b \prec c \prec a$, where $\prec$ represents the preference relation and a, b, c are outcomes. The reason given for such a constraint is the *money pump argument*, which goes like this: if $a \prec b$, the agent will pay to swap a for b; if $b \prec c$, the agent will pay to swap b for c; if $c \prec a$, the agent will pay to swap c for a; and so on—the agent can be made to give out money indefinitely.[15] **(3)** Conclude that the agent ought to satisfy constraints Y—and thus ought to have probabilistic credences. Thus we see that decision theoretic arguments too assume a telic account of justification: it is in virtue of how they relate to (pragmatic) aims that probabilistic credences are rational.

I have shown that the major arguments in the Bayesian literature are all telic.[16] Before examining the relevance of this fact for my purposes in this paper, let me explain how the telic/poric distinction relates to the literature on epistemic consequentialism. Berker (2013a), with whom this literature originates, defines what is now known as *epistemic consequentialism* to be the thesis that agents ought to believe what is most conducive to the aim of inquiry.[17] He writes: "according to this

[15] The money pump argument has roots in Ramsey (1926), and is later found in Davidson et al. (1955).

[16] Besides these, there are argument-types in Bayesianism that I have not discussed, including symmetry arguments (van Fraassen, 1989; Zabell, 2016), conservativeness arguments (Dietrich et al., 2016), and calibration arguments (van Fraassen, 1983; Lange, 1999). But I think it's fair to say that the arguments I have discussed are by far the most popular amongst contemporary Bayesians. Showing that they are telic and that there are good reasons for telics to be anti-subjectivists will suffice to account for the widespread rejection of subjectivism.

[17] Berker names this position *epistemic teleology*, but it has come to be known as *epistemic consequentialism*. Thus the two terms are often taken to be synonymous. But the point I am making in this passage is that consequentialism and what I am calling the *telic view* are not synonymous: the former is a special case of the latter. The terminology might be confusing, but unfortunately there are only so many words. I have tried to mitigate this by using "telic" and not, as he, "teleological".

picture, there are certain epistemic *ends* or *goals* that it is epistemically good for us to promote, and the question of what we *should* believe is determined by how well our believing conduces toward the fulfilling of those goals, or the furthering of those ends" (p. 340, emphasis original). Thus it is easy to see that epistemic consequentialism is a telic position: what matters when determining whether a credal state is justified is how it relates to the aim of inquiry.[18] However, there also exist telic positions besides epistemic consequentialism; in other words, not all telics are epistemic consequentialists. Indeed, epistemic consequentialists make two assumptions beyond telism. The first concerns the nature of the aim of epistemic attitudes: epistemic consequentialists assume that this aim is truth or accuracy. But if, as many Bayesians do, we take seriously the fact that credences are also that which rationalises action, we might consider pragmatic goodness to be an aim of credence, too—as indeed the proponents of Dutch Book and decision theoretic arguments do. And the second concerns the nature of the relationship to the aim of inquiry. As Berker notes, "what is distinctive about [epistemic consequentialism] is not just its taking value to be fundamental but moreover its attitude toward the nature of value and how we should respond to it. According to the [epistemic consequentialist], the proper response to value is to bring it about, and the proper response to disvalue is to stop it from being brought about: in short, for the [consequentialist] all value is 'to be promoted,' and all disvalue is 'to be prevented'" (p. 343). But it is possible to imagine a position whereby justification is a matter of relating to the aim of inquiry (a telic position), but where this relation is not one of conduciveness or promoting. For instance, Sylvan (2018, 2020) argues that justification is a matter of *respecting* the truth. Thus epistemic consequentialism is one among several possible telic views.

This is not to say however that the poric/telic disjunction exhausts the space of possible views about epistemic justification. Indeed, both the telic and the poric

[18] The accuracy arguments mentioned earlier in the section are widely acknowledged to be epistemically consequentalist, and this generates problems for them. For discussion, see Greaves (2013), Carr (2017), Pettigrew (2018) and Konek and Levinstein (2019).

epistemologist view epistemology as a means-ends endeavour of sorts: there is an epistemic good towards which the agent strives—a goal to be achieved on the basis of one's means: both take epistemology to be an activity oriented towards determining what is the case, an activity which agents perform from their limited perspective. I shall call this view a *means-ends* view of epistemology, and I develop it in more detail in §6.[19] But we can imagine accounts of epistemic justification that are not means-ends in this way. Take *evidentialism* for instance, the view originating with Feldman and Conee (1985, 2004) according to which an epistemic state is justified to the extent that it reflects the agent's evidence. Berker (2013b) argues that Conee and Feldman are not consequentalists: he says that "they do occasionally lapse into talk of 'epistemic value' and 'epistemic goals,' but I think this talk can be excised from their program without major loss" (p. 380). If he is right, their account of justification is not only non-consequentialist, but it is also non-telic, and non-means-aims. For them, epistemic justification is not a matter of truth-seeking or any other epistemic striving. I will discuss evidentialism in greater detail in §5, where I will consider various interpretations of evidentialism. But I want here to illustrate the point that there exist alternatives to means-ends epistemology.

Let us now circle back to my argument. As I announced earlier, I am going to argue that subjectivism is an appealing position on the poric approach. But if, as I suggested in this section, all the main Bayesians are telics, then should we not suspect that the widespread arguments against subjectivism presuppose a telic account of rationality? And if that is the case, would it not be natural to

[19] This label is imperfect. Indeed, it has an instrumental flavour, and evokes a view called *epistemic instrumentalism*, considered by Kelly (2003), according to which epistemic rationality is akin to practical (instrumental) rationality, in the sense that the agent determines her aim. On such an account of epistemic rationality, an agent whose aim it is to have false beliefs would be justified in believing that the earth is flat. But I want to stress that, the aims of means-ends epistemology are not particular to the agent or decidable by her; rather, the epistemic aims are those in relation to which epistemic states are evaluated, whether or not the agent additionally values these aims.

suspect that subjectivists and anti-subjectivists are engaged in a verbal dispute? That subjectivists and their critics do not disagree about anything substantive, but merely operate with different conceptions of justification? I will show in §4–5 that the first suspicion is correct: the over- and under-demandingness objections rely on telic assumptions, and crumble if these assumptions are replaced by poric ones. But I will also show that the second reaction would be too hasty. I will argue that we cannot do away with the poric approach. If this is right, it poses a dual problem for the status quo. Firstly, if, as I argue, porism and telism are incompatible and porism is indispensible, the standing of telism is seriously called into question. And secondly, if, as I argue, subjectivism is true on the poric approach, and the poric approach is indispensible, we will have to grapple with the fact that subjectivism is, in many ways, a terrible predicament.

3.

In this section, I present an argument for subjectivism, on the poric approach. I begin by making two (contentious!) assumptions, which I will discuss later on—dualism in §6, and idealisation in §4.

> *Dualism.* Agents have two means of inquiry: reason, and observation.
>
> *Idealisation.* Agents have their means of inquiry perfectly.

So, according to dualism, agents have two distinct ways of inquiring into the world: they have reason, on the one hand, and observation (or experience, or sense-perception) on the other. And according to idealisation, they have these means of inquiry perfectly. What does that mean? Philosophers have long discussed how successful our means of inquiry might be. For instance, epistemologists of perception have discussed whether and how observation could be a source of justified beliefs about the external world; that is, whether and how our capacity to observe could warrant the belief that, for instance, the book I am holding is red, rather than just the belief that the book seems red to me.[20] Similarly, epistemologists of

[20] This is the problem posed by Descartes (1993). See Lyons (2017) for an overview of the contemporary literature.

the *a priori* have debated whether and how reason could be a source of justified belief about non-empirical realms, such as the logical realm, the mathematical realm, the moral realm, and so on.[21] The idealisation assumption is designed to circumvent these challenges. It states (boldly!) that they have been answered, such that Bayesian agents are in a position to determine that the book really is red on the basis of observation, and that 2+2 really does equal 4 on the basis of reason.

It follows that taken together, the dualism and idealisation assumptions entail that Bayesian agents have two properties. Firstly, they are *perfect reasoners*. They are always in a position to recognise any proposition that can in principle be settled by reason, that is, any proposition decidable *a priori*, and always in a position to be certain of such propositions. This may include logical propositions, mathematical propositions, metaphysical propositions, moral propositions, and so on, depending on one's views on those realms of inquiry. And secondly, they are *perfect observers*. They are always in a position to recognise any proposition that can be settled on the basis of their observations, and always in a position to be certain of these propositions. So, for agents with such properties, propositions come in three kinds. The first kind are the propositions determinable *a priori*, those whose truth-value the agent can settle by exercising her capacity to reason. This leaves the propositions not decidable by reason, that is, the propositions determinable *a posteriori*, or empirical propositions. A further distinction between those is salient to our agent. On the one hand, there are the propositions she is in a position to settle on the basis of her observations; we can call those the propositions *about the observed*. And on the other, there are the propositions that she is not in a position to settle on the basis of her observations; we can call those the propositions *about the unobserved*.[22]

[21] See Russell (2014) for an overview of the literature.

[22] Some propositions may be determinable both *a priori* and *a posteriori*. For instance, if I want to know what 2+2 equals, I can calculate it mentally, or I can look at the result on my calculator. But this does not threaten my arguments. Indeed, because of the idealisation assumption, any proposition that is knowable *a priori* is already known by the agent. And as will become clearer, the propositions that are of interest to the poric

The next step is to model this agent with the help of the Bayesian mathematics. How can we model the *a priori*/*a posteriori* distinction? A very natural option is to use the trivial/non-trivial distinction. So, all propositions determinable *a priori* are modelled by Ω or $\emptyset$ (depending on whether they are true or false), and all the propositions not so determinable are represented by the non-trivial elements of $\mathcal{A}$. Among those, we can further distinguish between $E_1, E_2, ...$, the propositions about the observed, and $A_1, A_2, ...$, the other propositions, about the unobserved. We now find ourselves with an interpretation of the Bayesian algebra of entertained propositions. This interpretation has, I think, a distinctly Bayesian flavour.[23] The elements of $\mathcal{A}$ are propositions that the agent might come to learn by empirical inquiry: propositions that might constitute, or come to constitute evidence—that might become E_is. By contrast, the trivial propositions Ω and $\emptyset$ are those whose truth-value cannot be determined by looking in the world. Note that on this interpretation, trivial omniscience is not mere logical omniscience: it is omniscience about all *a priori* matters.

Now, onto the argument for subjectivism. Remember, subjectivism is the conjunction of the necessity claim, which says that agents ought to be trivially and evidentially omniscient, and the sufficiency claim, which says that there are no other requirements of rationality. Here goes the argument.

1. An agent is always in a position to be certain of the truth-value of *a priori* propositions. (Dualism and idealisation assumptions.)
2. The *a priori* propositions are Ω and $\emptyset$. (Modelling assumption.)
3. An agent is always in a position to be certain that Ω is true and $\emptyset$ is false. (1, 2)
4. An agent's credences should be such that $p(\Omega) = 1$ and $p(\emptyset) = 0$. (Poric justification, 3.)

5. An agent is always in a position to be certain of the truth-value of propositions about the observed. (Dualism and idealisation assumptions.)

---

are those whose truth-value the agent's means do not suffice to determine.

[23] Note too that it is a much better candidate than the orthodox metaphysical interpretation to account for uncertainty about the necessary *a posteriori*. See Chalmers (2011).

6. The propositions about the observed are $E_1, E_2, ...$ (Modelling assumption.)

7. An agent is always in a position to be certain that E_i is true and $\neg E_i$ is false, for all i. (5, 6.)

8. An agent's credences should be such that $p(E_i) = 1$ and $p(E_i) = 0$, for all i. (Poric justification, 7.)

9. *Necessity claim.* Agents ought to be trivially and evidentially omniscient. (4, 8.)

10. Neither reason nor observation warrants particular credences in propositions about the unobserved.

11. The agent has no means of inquiry beyond reason and observation. (Dualism assumption.)

12. *Sufficiency claim.* There are no other requirements of rationality. (Poric justification, 10, 11.)

13. *Subjectivism.* (9, 12.)

So why be a subjectivist? Because trivial and evidential omniscience, but nothing else, is warranted by Bayesian agents' means of inquiry. These agents' perfect capacity to reason and to observe warrant certainty in the *a priori* propositions, and the propositions about the observed. But these capacities do not warrant any particular epistemic attitude towards the third kind of propositions under consideration: the propositions about the unobserved. As such, there is nothing that agents ought to believe about them.

4.

We have made some progress towards subjectivism: we have an argument in favour of the position. But it remains to be shown that its proponent (along the lines above) can address the objections to the necessity and the sufficiency claims. I will begin with the over-demandingness objection in this section, and will move on the under-demandingness objection in the next. For each, I will concede that the telic epistemologist should reject the relevant claim in light of the argument, but argue

that the poric epistemologist should not. The first part of this argument—that the telic Bayesian falls prey to these objections—will play a dual dialectical role. On the one hand, and given my argument in §2 that all major Bayesian arguments are telic, it will explain why these objections have been so successful in persuading Bayesians away from subjectivism. On the other hand, it will function as a challenge to the status quo in Bayesian epistemology. For if subjectivism is false on the telic view, but it is true on the poric view with which—I will argue—we cannot do away, we might be forced to reject telism altogether. I discuss this in §6. In the meantime, my aim is not so much to entrench the telic's rift from subjectivism as it is to show that there is no such rift for the poric. Accordingly, the standard of argument will be higher for the latter goal than for the former: I will do little more than rehearse already widely accepted explanations for the falsehood of subjectivism on the telic approach, but I will seek to actively convince the reader that the objections fail on the poric approach.

With the dialectic made clear, we are ready to begin. Let me first reproduce the over-demandingness objection.

Over-demandingness objection.

1. Agents ought to be trivially and evidentially omniscient. (Necessity Claim)

2. Ought implies can. (Premise)

3. Agents can be trivially and evidentially omniscient. (1, 2)

4. They cannot. (Premise)

5. So, the necessity claim is false.

Let us begin with premise **2**. There are two kinds of telics in the literature. The first kind seeks to issue *guidance* to the agent on what to believe; for instance, Bradley seeks to construct a theory which "would provide guidance on how bounded agents should represent the uncertainty they face, how they should revise their opinions as a result of experience" (p. xiii, 2017). The second kind seeks to evaluate or *appraise* the agent's epistemic state; Pettigrew for instance wants to know "when it

is appropriate to criticize an individual for their logical ignorance" (forthcoming).[24] How does premise **2** fare for each type of telic Bayesian?

It is straightforward that Bayesians who seek to guide cannot reject it: were it not the case that ought implies can, judgments of rationality could not serve as advice. (What kind of advice is unfollowable advice?) It is less straightforward, but still true I think, that Bayesians who seek to appraise cannot either. There are two things that one can do under the guise of appraising: (i) comparing an agent's epistemic state to the perfect epistemic state; and (ii) comparing an agent's epistemic state to the best such state that the agent could have been in. If telic Bayesians were in the business of doing (i), then they would be in a position to reject premise **2**. But they would also need to mandate empirical omniscience—a principle according to which agents must be certain of the truth-value of all empirical propositions. This is because the perfect epistemic state is that whereby all truths are believed with certainty and all falsehoods correspondingly disbelieved. Given that they do not mandate empirical omniscience, we can infer that it is (ii) in which they are involved. Thus we should interpret Pettigrew's question of when it is "appropriate" to criticise someone's beliefs, as asking which of the beliefs that the agent could have had (for some sense of "could") was the right one to have had. For the appraising Bayesian too, ought implies can. (I should remark that the sense of "can" plausibly differs in the guiding and appraising projects[25]—but it does not differ in a way that is relevant to us; namely, in whether Bayesian agents *can* be trivially and evidentially omniscient.) Therefore, telic Bayesians of all stripes must

[24] Pettigrew steers clear of the guidance aim in response to the recognition that human agents do not in general have control over what they believe: *epistemic voluntarism* is false. While I agree with him that epistemic voluntarism is false in general, I think that the cases in which (telic) Bayesian epistemology is most useful are cases where it is true, and I suspect that those are the cases that Bradley has in mind. For instance, we can easily imagine a scientist genuinely wondering what credences to form in response to a body of evidence.

[25] For instance, guiding Bayesians must interpret "can" in such a way that it implies epistemic voluntarism (see fn. 24), whereas appraising Bayesians need not.

accept premise **2**.

Since it is because of the telic's intention to guide and/or appraise that he must accept premise **2**, and since the poric Bayesian need not share that aim, might it be that she can reject premise **2**? No. For her, an epistemic state is justified just in case it is warranted by the agent's means of inquiry. In other words, what it means to say that an agent ought to have a particular epistemic state is that her means of inquiry allow her to have it, warrantedly. So the poric Bayesian too must accept premise **2**: it is constitutive of what she means by rationality that ought implies can (in the sense of "can" relevant to the over-demandingness objection).

Let us now turn our attention to premise **4**. Assuming that we take trivial omniscience to entail logical omniscience, and that we adopt an externalist account of evidence, it is clear, as we saw in §1, that if we take Bayesian agents to be (models of) real human agents, this premise is true. Actual humans are not always in a position to recognise logical and evidential propositions as such. But Bayesians are well known to be engaged in idealised epistemology: epistemology that brushes over humans' various cognitive and perceptual limitations. And if an idealisation assumption can be justified on which Bayesian agents can be trivially and evidentially omniscient, the necessity claim might be protected from the over-demandingness objection. For instance, on the poric position outlined in §3, the theorist is not particularly interested in real human agents. Instead, the agents that she considers are highly idealised. In fact, they are perfect reasoners and perfect observers, who are, by assumption, always in a position to settle any proposition determinable *a priori*, and always in a position to settle any external-world proposition about the observed. It follows that, on this view, agents can be trivially and evidentially omniscient, and premise **4** holds. So, since it is with the idealisation assumption, that, perhaps unsurprisingly, the necessity claim stands or falls, let us examine how it fares on both the telic approach and the poric approach.

Can telics justifiably accept the idealisation assumption? It is noteworthy that few think they can: Bayesians have long believed idealisation assumptions to be ultimately unwarranted, and the recent trend towards de-idealisation is gaining

momentum. And indeed, this follows from the considerations rehearsed above. If telics seek to apparise and/or guide real, human agents, any idealisation assumption will translate into an unachievable requirements on their agents, and so, ultimately, in the failure of their project. In light of this, the role of idealisation for them is transitional at best—that is, if they do idealise, it is with the intention of subsequently weakening the idealisation assumptions.[26] The goals of telic Bayesianism undermine idealisation assumptions.

But what about poric Bayesianism? We saw in §2 that the key question on the poric approach is that of which epistemic states an agent might achieve on the basis of her means. And we saw in §3 that a particular domain of inquiry stands out: the empirical unobserved. Indeed, we saw that the dualism assumption divides propositions into three types: those that reason might settle (the *a priori* propositions), those that observation might settle (the propositions about the observed), and those that neither reason nor (heretofore) observation suffices to settle—propositions about the unobserved. And it is this latter type, which is the proper target of poric Bayesianism. Indeed, all those who use Bayesian methods in epistemology are interested in empirical questions, which go beyond the agent's evidence. And this is unsurprising: the Bayesian formalism is not a particularly promising method of epistemological inquiry into non-empirical realms.[27] Bayesianism decidedly concerns the empirical.

And, what credences in propositions about the empirical unobserved are warranted by our means of inquiry? The argument for subjectivism presented in the previous section suggests that the answer to this question is: none. Our means of inquiry do not warrant any particular credences in propositions about the unobserved. For the poric Bayesian, and in accordance with widespread intuitions,

[26] For a discussion of transitional idealisation, see Staffel (2017) and de Bona and Staffel (2018).

[27] As far as I am aware, no one has attempted to use Bayesian methods in mathematical or metaphysical epistemology. There have been some attempts to use Bayesian methods in moral epistemology, for instance by MacAskill and Ord (2020), but many problems arise and solving them requires significant departures from the orthodox Bayesian framework.

this is a terrible conclusion: there is a domain of inquiry for which our means do not suffice to support any particular epistemic attitude. We are now in a position to appreciate the dialectical role of the idealisation assumption. By assuming that agents have their capacities of inquiry perfectly, the poric theorist is *conceding* more to agents than they in fact have. For if *even* agents who have these means of inquiry perfectly cannot have justified beliefs about the unobserved, agents who have them imperfectly certainly cannot.

So the poric Bayesian makes the idealisation assumption and thus rejects premise **4**, not because she is naive or confused about the capacities of real human agents, but because she wishes to concede to such agents more than they have, in order to determine what they might be able to achieve, epistemically. This is in sharp contrast with the situation of the telic Bayesian, for whom the idealisation assumption is not a concession but a injunction. Once we get the broader aims of telic and poric Bayesians plainly into view, it becomes clear that they should respond differently to the over-demandingness objection. The telic Bayesian indeed has no choice but to reject the necessity claim, for it makes an impossible demand to his agents. But the poric Bayesian should embrace the necessity claim, because for her, it plays the role not of a demand but of a concession: it gives to the agent more than she has, so as to see how far she could in principle go.

5.

This brings us to the other widespread objection to subjectivism: the under-demandingness objection. Just like with the over-demandingness objection in the previous section, my aim in this section is to show that, although there are good reasons to take the under-demandingness objection to be decisive against subjectivism on the telic approach, this is not so on the poric approach to rationality. Thus I will have accounted for the widespread endorsement of the under-demandingness objection to the sufficiency claim, shown that that the objection does not threaten the poric case for subjectivism, and confirmed that there is a genuine tension between the telic and poric approaches. Let me start by reproducing the argu-

ment.

Under-demandingness objection.

1. All that's rationally required of an agent is for her to be certain in the truth of trivial and evidential propositions. (Sufficiency Claim)

2. A credence of .99 in the proposition that an unknown coin will land heads is rational. (1)

3. It is not. (Premise)

4. So, the sufficiency claim is false.

We should begin by noting that this argument is less sophisticated than the over-demandingness objection. Its structure is very basic: we consider a direct consequence of the sufficiency claim, and rule on the basis of intuition that it is false. Now, a mere hunch can clearly not be reason enough to reject the sufficiency claim. We need (at least!) to explain why we have this intuition, in a way that exposes some philosophical consideration. Bayesians have attempted this in three ways—in what follows, I review each of them in turn.

5.1.

The thought at the heart of the first is that it is irrational for an agent to have a very high credence in the proposition that a coin will land heads because such a credence fails to match up with the objective chance that it will. We can call this version of the argument the *chance-based version.* So far in this paper, we have used probability functions to represent epistemic phenomena: agents' degrees of belief. But these functions are also used to represent ontic, or mind-independent phenomena; so interpreted, they are known as *chances.*[28] And a popular thought is that credences and chances are not unrelated—that the credences we ought to have are constrained in some way by ontic probabilities. To capture this thought, philosophers have formulated so-called *bridge principles*, the most famous of which is Lewis' Principal Principle (1980). Which formulation of a bridge principle is

[28] For an overview of different ontic interpretations of probability, see Hájek (2019).

the right one is a surprisingly complex and subtle matter, but the details are not important here.[29] What is important—what is at the core of all contestant bridge principles—is that an agent's credence p in a chancy proposition X, given that the correct chance function is ch, should be equal to the value that ch ascribes to X. So, where C_{ch} denotes the proposition that the correct chance function is ch:

$$p(X|C_{ch}) = ch(X)$$

Pettigrew has proposed a Dutch Book argument (2020) and several accuracy arguments (2012, 2013a, 2016a) for bridge principles. (As we saw in §2, these argument-types are telic.) But assuming that at least one of these arguments succeeds, and thus that a bridge principle indeed holds, it still does not follow that premise **3** is true—that an agent with a very high credence that the coin will land heads is irrational. For what bridge principles express is that there ought to be a particular relation between the agent's credence in the chancy proposition X and the proposition that the correct chances are given by ch. But if the agent does not know which chance function is the correct one, a bridge principle alone cannot constrain her credences. This is often expressed by saying that bridge principles are mere coherence requirements—or, in my terminology from §1, they are formal norms. It follows that, for it to be possible to run the chance-based version of the under-demandingness objection on the poric approach, it must be the case that a Bayesian agent is in a position to become certain of which is the correct chance function, on the basis of her means of inquiry—and it is not.

Accounts of chance can be divided in two kinds.[30] On the *non-reductionist* accounts, such as the hypothetical frequency or the propensity accounts, chances

[29] Pettigrew gives an illuminating presentation and critical discussion of numerous versions (2016a, pt. II).

[30] I have defined the reductionist/non-reductionist distinction, not in the typical way to mean ir/reducible to *non-modal* facts, but to mean ir/reducible to facts about the *observed*. So, on my way of drawing the distinction, it is plausible to categorise the best-systems account as non-reductionist, for although it is reducible to non-modal facts, it is not reducible to observed facts—one cannot have observed all facts past and future.

are not reducible to facts about the observed. On the hypothetical frequency account for instance, the chance of a coin landing heads is the limiting frequency with which it would land heads, were it flipped an infinite number of times. It is plain that chances on these accounts cannot be determined by observation.[31] By contrast, on the *reductionist* accounts, such as the finite frequency account, chances are a kind of summary of observed facts. For instance, on the finite frequency account, the chance of a coin which has been tossed ten times landing heads is the frequency with which it has actually landed heads. Chances on these accounts can be determined by observation.

The non-reductionist accounts of chance cannot help construct a poric version of the under-demandingness objection, since, by definition, these chances cannot be determined by observation. But, can we not conjoin a bridge principle with a reductionist account of chance to support premise **3**? We cannot. This is because, in order for the agent to be in a position to determine the (reductionistically interpreted) chance of a proposition, it must be for that proposition to be about the observed. But in order to undermine the sufficiency claim, it has to be the case that the proposition considered is one about the unobserved—one not determinable on the basis of observation. Indeed, the anti-subjectivist insists that there are propositions beyond the observed in which agents can have justified credences. It follows that, whether on a non-reductionist or a reductionist account of chance, there cannot be a poric chance-based version of the under-demandingness objection.

5.2.

This brings us to the second widespread reason for rejecting the sufficiency claim, which I take to be the most widely endorsed. According to this second group

[31] Although, some authors have attempted to ground chanciness in empirical features of the world, by appealing to what is known as the *method of arbitrary functions*. See for example Strevens (2013). For a counter-argument, see "Objectivity and the Method of Arbitrary Functions", published as Canson (forthcoming).

of Bayesians, premise **3** is true—it is irrational for an agent to have such a high credence that the coin will land heads—because, in the absence of evidence, having such a wild credence is taking an undue risk. We can call this version of the under-demandingness objection the *risk-based version.*

This version of the objection consists in a risk-based argument to the effect that an agent should have minimally opinionated credences—she should satisfy a *principle of indifference.* Such principles state that, in the absence of evidence, an agent's credence function should be flat; that is, agents should assign equal credences to (inconsistent) propositions. An indifference principle was first formulated by Bernoulli (1837/1954) and comes to Bayesians through Keynes (1921) and Carnap (1950); adherence to such a principle is the hallmark of what is known as *objective Bayesianism.*[32]

There are two risk-based arguments for the principle of indifference. The first, by Williamson (2010), goes as follows. An agent with non-indifferent credences will, because of the constitutive relationship between credences and action, be disposed to act in very risky ways. For instance, the person with a very high credence that the (unknown) coin will land heads will be disposed to bet a significant portion of his wealth on that outcome. But doing so, the argument goes, goes against what it would be best for the agent to do, and so the agent ought to have a milder opinion. The second risk-based argument, by Pettigrew (2016a,b), has a similar high-level structure (though it differs significantly in the details). An agent with a very high credence that the (unknown) coin will land heads takes a big epistemic risk that the agent with a middling credence in that proposition does not—she risks being very

[32] Objective Bayesianism is meant to be the opposite of what is sometimes known as Subjective Bayesianism, and which I have called subjectivism. But as I have defined it, subjectivism is the conjunction of two claims, the necessity claim and the sufficiency claim, and it is only the latter that objective Bayesians reject. Furthermore, they reject it by appealing to an indifference principle, but as we have just seen, there are more than just one way to reject it. We can say though that the so-called objective Bayesians are the most objective of Bayesians in one sense: they impose the most stringent constraints on all of an agent's credences.

inaccurate. And, the argument goes, agents ought not take undue alethic risks: they ought to satisfy the principle of indifference.

Let us see how these risk-based arguments for the principle of indifference feature in a response to the sufficiency claim. Firstly, they vindicate premise **3** of our argument above: the agent is faced with two propositions, heads and tails, and the principle of indifference entails that she ought to have equal credence in both. Moreover, they are typically thought to vindicate intuitions of rationality more broadly. For consider an agent who has a low credence that the sun will rise tomorrow. Intuitively, this agent is irrational. Risk-based proponents of the principle of indifference explain this intuition in the following way. Before he had any evidence about the sun's movements, the rational credence function for him to have had would have been the indifferent one: this follows from the indifference principle. Since then, he has received plenty of evidence to the effect that the sun has risen every day. Given the principle of conditionalisation (and given some constraints on conditional credences), his credence that the sun will rise tomorrow should be greater than half.[33] It follows, let us grant, that risk-based arguments refute the sufficiency claim. But note that both of these arguments are telic: it is because not having indifferent credences risks being (pragmatically or alethically) bad for the agent that she ought to satisfy the indifference principle. Furthermore, it is hard to see how this telic argument could be reformulated in a poric way: it relies unavoidably on the expected goodness of adopting particular credences. As such, to the extent that a risk-based version of the objection is successful against the sufficiency claim, it is so in virtue of being telic. While it may constitute a good reason for telic Bayesians to reject subjectivism, it does not threaten poric Bayesians' adherence to the position.

5.3.

This brings us to the third version of the under-demandingness objection. According to this third group of Bayesians, premise **3** is true—it is irrational for an

[33] I argue against this line of reasoning in "Bayesianism and the Problem of Induction".

agent to have such a high credence that the coin will land heads—because such an epistemic attitude fails to reflect the agent's evidence. Proponents of this view often cite Hume as asserting that "the wise man [...] proportions his beliefs to the evidence" (1748, 10.1–4). We can call this version of the argument the *evidentialist version.*

Evidentialists in the Bayesian tradition are particularly noted for their defence of the principle of indifference, which takes place in two steps. They start from the consideration that, since our agent does not have any evidence about the outcome of the coin flip, her evidence supports both outcomes—heads and tails—equally.[34] They continue with the contention that, as Keynes (1921) puts it, "if the evidence affords no ground for attributing unequal probabilities to the alternative predictions, it seems to follow that they must be equal" (p. 45). Or as White (2009) more recently puts it, agents' credences should "reflect their evidence (or lack of it)" (p. 171). Thus in the absence of evidence, agents should not assign a higher credence to one proposition over another. But evidentialist arguments can also be given beyond situations where the agent lacks evidence. For instance, an evidentialist might claim that our evidence that the sun has risen every day so far supports a high degree of confidence that it will rise again tomorrow, and conclude that an agent ought to have a high credence that it will.[35] This means that, if successful, these arguments have a compelling anti-subjectivist potential: they can be used to vindicate credal constraints in a wide variety of scenarios.

In general, evidentialism can be characterised as the conjunction of two claims. **(1)** According to the descriptive claim, for any body of evidence, there is a unique credence function that is supported by that evidence.[36] **(2)** According to the norm-

[34] I should note that this claim does not seem right to me; rather, I would be tempted to say that there is no particular credence that lack of evidence supports, which would be more in line with the poric spirit of this paper. But this will not be relevant for our purposes, as we will see.

[35] Thus evidentialists can avoid appealing to the credences the agent had before receiving evidence in order to determine what credences it is rational for them to have at a later time.

[36] Some think that it is not a unique credence function but a unique *representor* (a rep-

ative claim, an agent ought to have the credences supported by her evidence. (Alternatively, her credences ought to "reflect" her evidence, in White's terminology (2009); or they ought to "respect" it, in Sliwa and Horowitz' words (2015).) In the rest of this subsection, I will examine whether evidentialism threatens my contention that the sufficiency claim is true on the poric approach. I will consider two interpretations of evidentialism. The first one, already mentioned in §2, is orthogonal to the means-ends conception of epistemology that I am working within, and as such, irrelevant to my purposes. The second one can be situated within means-ends epistemology, but I will argue does not threaten my contention.

Let us begin with the first interpretation of evidentialism. On this interpretation, it is *constitutive* of rationality that an agent's credences reflect the evidence. To deform Hume's dictum: to be wise just is to proportion one's beliefs to the evidence. On this interpretation, the normative claim **(2)** is analytic, and all the work is done by the descriptive claim **(1)**, which, it should be emphasised, has long been taken to be the Achilles heel of the position. Indeed, Keynes and his successors were widely criticised for failing to provide an account of evidential support, despite their views relying on the existence of one. Nonetheless, since many contemporary Bayesians—including Sliwa and Horowitz (2015), Schoenfield (2015, 2017b), or Dorst (forthcoming)—use the notion of evidential support, I will assume that whether an account of it can be formulated is still a live question. What is interesting, for my purposes, about this interpretation of evidentialism, is that it is neither telic nor poric; it in fact refuses what I called in §2 a means-ends conception of epistemology, whereby inquiring agents have a goal that they ought to achieve. Indeed in that section, I quoted Berker as claiming of Conee and Feld-

resentor is a set of credence functions) that is supported by the evidence—they are the so-called *imprecise Bayesians*. The debate between precise and imprecise Bayesians concerns the level of grain of the epistemic state supported by the evidence; but what matters here is the assumption that there is a unique epistemic state which is indeed supported by the evidence. So I leave aside, here, and in fact elsewhere in this paper, concerns about the grain of epistemic attitudes. See Levi (1985) and Joyce (2005) for early texts, and Bradley (2019) for an overview.

man, the most prominent evidentialists outside the Bayesian tradition, that "talk [of epistemic goals] can be excised from their program without major loss" (2013a, p. 380). Instead, for this kind of evidentialist, the epistemic agent is not attempting to determine what is the case, she is merely mirroring her evidence. Rationality on the constitutive interpretation is just a matter of calibrating one's credences to the evidence, and not a matter of seeking the truth. As Berker (2013b) puts it, it is "sideways-looking" (p. 377). As such, this interpretation is orthogonal to the framework within which my arguments in this paper operate.

Let us then turn to a second interpretation of evidentialism, consistent with the means-ends conception. How can evidentialism be formulated within such a framework?[37] A natural answer is that agents' beliefs ought to mirror their evidence because this is a good way of achieving the ultimate epistemic aim, namely, determining what is the case. So, on this second interpretation of evidentialism, believing in line with one's evidence is a means to achieving epistemic aims.[38] Note that this has a distinctly poric flavour. Might it be that evidentialism, so interpreted, yields a poric case against the sufficiency claim?

To answer this question, let us use a case proposed by Sliwa and Horowitz: "Anton is an anesthesiologist, trying to determine which dosage of pain medication is best for his patient: A or B. To figure this out, Anton assesses some fairly complex medical evidence. When evaluated correctly, this kind of evidence determ-

[37] For variants of this question, see Easwaran and Fitelson (2012), Schoenfield (2015) and Easwaran (2017).

[38] Pettigrew (2013b) writes, in relation to evidentialism, that "accuracy is not the only goal of credences: there is also the goal of matching one's credences to one's evidence" (p. 579). He goes on to consider how the goals might relate, and concludes that the only viable option if one is to remain within what I have called the means-ends framework is to hold that the "evidential goal is not an independent goal at all, but rather a byproduct of the goal of accuracy" (p. 579). There is a literal reading of Pettigrew's claims on which he disagrees with me: he takes reflecting the evidence to be a goal, albeit subsidiary to that of accuracy, whereas I take it to be a means of achieving accuracy. But I think there is less disagreement than it might appear: to be a subsidiary goal to X, and to be a means to X seem very similar things to me.

ines which dose is right for the patient. After thinking hard about the evidence, Anton becomes highly confident that dose B is right. In fact, Anton has reasoned correctly; his evidence strongly supports that B is the correct dose" (p. 2836).[39] To fit this case in the two-step evidentialist recipe I provided above: **(1)** Anton's evidence supports a high credence that B is the correct dose; and **(2)** he should therefore adopt a high adopt a high credence in B.

Now, suppose that the descriptive claim **(1)** is indeed true; that is, that there is a (good) notion of evidential support such that Anton's evidence supports a high credence in B. Then, provided that **(2)** holds, it follows that the sufficiency claim is false: there are requirements of rationality beyond those maintained by subjectivists. But this interpretation of evidentialism is not poric. For if we interpret the notion of evidential support from the descriptive claim as one of poric warrant—that is, as one warranted by the agent's means of inquiry: observation and reason—it cannot be the case that Anton's evidence warrants a high credence in B. Indeed, whether B is an effective treatment is not observed, and it is not determinable *a priori*. Thus the claim that Anton's evidence supports a high credence in B must operate with a non-poric conception of evidential support.[40] To conclude: if evidentialism underpins the under-demandingness objection, it is not poric; or contrapositionally, if evidentialism is poric, it does not underpin the under-demandingness objection.

Let me conclude on this section. The success of the under-demandingness objection relies on there being a reason to accept a claim like premise **3**: that it is irrational to have a very high credence that a coin about which nothing is known will land heads. In this section, I examined three versions of the objection, each putting forward a different reason for **3**. The chance-based version failed for being

[39] There is nothing special about this case, it is fairly standard within the formal evidentialist literature.

[40] On the poric interpretation of evidentialism, although **(1)** is false, we get **(2)** for free—to be justified for the poric just means to be warranted by the means of inquiry. But on the non-poric approach, the normative claim is non-analytic and must be defended: Why accept that believing in line with one's evidence is a good way of achieving accuracy?

a mere coherence requirement. The risk-based version was more convincing, but was distinctly telic. And the evidentialist version was either non-means-ends (thus not relevant to us), non-poric, or non-successful. I started this section by referencing the strong and widespread intuition to the effect that the agent of premise **3** is indeed irrational, and by advocating for a philosophical account of this intuition. We are now in a position to provide one. For the telic Bayesian, the intuition that the agent is irrational expresses a judgement, maybe to the effect that she would be badly off if she had this credence. But for the poric Bayesian, it expresses a kind of lament—how could the agent's means of inquiry warrant so little?

6.

My primary aim in this paper has been to provide an answer to the titular question: why subjectivism? The answer I have provided is that, on the poric account of rationality, according to which epistemic states are justified to the extent that they are warranted by the agent's means of inquiry, subjectivism is an appealing view. Not only is there a convincing argument in favour of the position (§3), but the typical arguments against it fail to have the required grip (§4–5). But, alongside this, I have also tried, if not to vindicate, at least to countenance the widespread anti-subjectivist beliefs, by conceding that they might be appealing on the alternative (and widely assumed), telic account of rationality. It follows that the debate between subjectivists and anti-subjectivists is best seen, not as a debate about particular putative norms of rationality, but as a debate about which approach to epistemic justification to adopt.

A natural reaction to this state of affairs would be to rule the telic/poric dispute—and the subjectivist/anti-subjectivist dispute with it—as merely verbal. But I think we should resist this urge. In general, in order to ascertain whether two opponents are talking to or past one another, one must determine whether they are trying to answer the same question. In the case of the poric and the telic, it is plausible that they are. Firstly, note that for both, epistemology is a goal-directed endeavour. This is obvious in the case of the telic, but also holds of the poric.

Indeed, the poric mandate to adopt the epistemic states warranted by one's means of inquiry cannot even be formulated without presupposing the existence of an epistemic aim: what is warranted is what can be built *up* from the means—built that is, in the direction of the aim. (There is an interesting question, of course, about what it is that is to be achieved—whether the goal is alethic, or pragmatic, or both—what shape it takes beyond this—but what matters here is a more abstract question about the structure of epistemic endeavours.) And secondly, note that achievement is a matter, not just of where one ends up, but also of what one started with. That the poric presupposes this is obvious, but the telic does too. Indeed the typical telic is not merely instructing agents to conform to the epistemic goal(s): he does not rule for instance that agents ought to have all and only true beliefs. Instead, he insists that agents ought to have the best beliefs that they could have, given their epistemic standing. In sum, the telic and the poric can plausibly be portrayed as trying to answer the same question: that of what agents can epistemically achieve on the basis of their means. They both operate within what I have called *means-ends* epistemology. And given that they are both trying to answer the same question, but that their responses differ, at most one of them can be right. This is why I think that a diagnosis of verbal dispute would be misguided.[41]

The telic could resist this, most obviously by denying that he sees epistemic endeavours through the lens of achievement; or in other words, that he is not engaged in means-ends epistemology. He would then, of course, have to explain what he thinks epistemic activity is about. But this insistence that the subjectivism/anti-subjectivism debate is merely verbal would not frustrate my main aim. Indeed, if I am to convince the reader to take subjectivism seriously, I need only to persuade her of two weaker claims: **(1)** that the means-ends conception of epistemic activity is a valuable one; and **(2)** that the poric account of epistemic justification plays a

[41] By contrast, the evidentialist who is involved in what Berker (2013b) calls a "sideways" endeavour, namely that of merely reflecting one's evidence, is not asking about what can be achieved. So were this evidentialist and the typical telic Bayesian engaged in a dispute, it would be a verbal one.

valuable role on this conception. And I think that these claims are clearly true. An important part of both scientific and quotidian endeavours is finding out about the world. And the way that we go about this kind of inquiry is using the means at our disposal: reason, the senses, etc. So whether one accepts my claim that the telic and the poric are actually disagreeing, one must accept my claim that poric epistemology is important—and thus, must take subjectivism seriously.

This is, in some sense, unfortunate. For subjectivism is a very pessimistic answer to the question of what one can epistemically achieve: if what I have said in this paper is right, then there are no epistemic attitudes that are warranted towards propositions in the relevant domain of inquiry, the empirical unobserved. But note that I have left one assumption unexamined: the dualism assumption, according to which agents have two means of inquiry—reason, and observation. In the rest of this paper, I examine various ways of rejecting the dualism assumption, and I conclude pessimistically.

There are three ways in which one could reject the dualism assumption. **(a)** One could argue that agents lack one of the two means posited by the assumption; agents lack either observation or reason as a means of inquiry. Now, not only will this strategy clearly not help us escape the grip of subjectivism, but it would also force us to abandon the building blocks of Bayesianism: agents must be able to reason if they are to satisfy probabilism, and to observe if they are to satisfy conditionalisation. This brings us to: **(b)** one could reject that observation and reason are two distinct capacities. Once again, this strategy will clearly not go very far in helping us refute subjectivism, and once again, it would force us to abandon some foundations of Bayesian epistemology; in particular, the contention that learning takes place through the acquisition of (pure) observational evidence.[42] This in turn brings us to: **(c)** one could reject that observation and reason are the only means of inquiry. This third option seems highly promising, for there is a widespread way by which people come to have justified beliefs about the world

[42] The claim that observation is thought-laden is clearly plausible—and I think constitutes an under-explored challenge to Bayesian methodology. But this is for another time.

beyond reason and observation: testimony.

This has been extensively discussed in the Bayesian literature, under the banner of *deference principles*. The details are subtle, but we can gloss over them for our purposes. A deference principle says that an agent's credence p in a proposition X, given that an expert's credence function is p^+, should be equal to the value that p^+ ascribes to X. So, where E_{p^+} denotes the proposition that the expert's credence function is p^+:

$$p(X|E_{p^+}) = p^+(X)$$

Where the expert is one's better informed future self, we call this a *reflection principle* (van Fraassen, 1984), and where the expert is a better informed other agent, we call it an *expert deference principle* (Elga, 2007). Various arguments have been proposed, including a Dutch Book argument by van Fraassen (1984), and an accuracy argument by Easwaran (2013). Now, like chance functions (as discussed in §5.1), deference principles can only help constrain an agent's credences if she in fact comes to know the shape of expert functions. This is where testimony might come in. But in fact, it might not need to: expert functions are possible to come to know by reason—reasoning about what one will believe in cases of reflection—and observation—hearing from experts about their credences in cases of expert deference. One might interpret this as counting testimony as a kind of observation: the observation of others' credences.

But whether one sees testimony as a kind of observation or as an independent means of inquiry, one is unlikely to get very far. For if Bayesian epistemology applies to everyone—as it surely does if it applies to anyone—then that which we can come to know by testimony is limited to what others can come to know by observation. And many—most!—of the propositions of interest to an agent are not only unobserved by herself, but by all agents. (Will this coin land heads? Will the sun rise tomorrow? Will my loved ones become infected with Covid-19?) This is so in large part because credences guide action, and the propositions that matter to the choice of the deliberating agent are about the future. When the rain-averse agent wonders whether he ought to take his umbrella, what matters to him

is whether it will rain after he leaves the house. When the policy-maker deliberates on climate action, what matters to her is what will happen at the time she plans for the policies to be implemented. And what additional means of inquiry that agents have could bridge that gap? So I conclude, pessimistically, that subjectivism is unfortunately true.

2. Bayesianism and the Problem of Induction

Abstract. It is widely held that Bayesian epistemology and the problem of induction are closely related. However, there is no consensus on the nature of this relation. In this paper, I examine and reject the two widespread views on this issue. I then present my own view, on which the problem of induction resurfaces intact in Bayesian epistemology, in the form of the question of which credences agents ought to have in propositions about the unobserved. I show that this has several revisionary consequences, concerning the role of conditionalisation in addressing the problem of induction, the need for superbabies, the relevance of the priors, and the standing of objectivism and subjectivism.

0.

Many take Bayesian epistemology to be relevant to the problem of induction in some way.[1] In fact, Bayesianism is often presented, in courses and encyclopaedia entries, as making significant headway towards solving the problem. More than that, it is sometimes claimed that it is because of its success vis-à-vis induction that Bayesianism is so popular. But agreement breaks down when we ask what precisely the problem is, and how Bayesianism addresses it. My aim in this paper is to reconstruct and refute the two widely-held views on this matter, and to build on my arguments about these two views to propose my own view, on which the problem

[1] In this paper, I am solely concerned with the traditional problem of induction usually associated with Hume, and not with Goodman's new riddle (1983).

of induction resurfaces intact in Bayesian epistemology. Four claims of interest follow from my arguments. The first is that, contrary to what the proponents of both existing views contend, a successful defence of conditionalisation does not constitute progress toward a solution to Hume's problem. The second is that the contentious appeal to superbabies need not be made. The third is that a widespread strategy to avoid engaging with the problem of induction—namely, insisting that all depends on the problem of the priors, as if it was a separate problem—is infelicitous; the problem cannot be pushed back thusly. The fourth is that my analysis positions subjectivism, objectivism, and other putative alternatives as direct responses to the problem of induction.

Let us begin by outlining the structure of the problem of induction. Firstly, there is a goal of induction, that which we are trying to establish. (As I will show, the two Bayesian camps disagree on what this goal is.) Secondly, there is a problem with achieving this goal. This has two components: an importance component—it is important for the goal to be achieved—and a difficulty component—it is not immediately clear or obvious how this goal is to be achieved. And finally, there are ways of addressing the problem. One might provide a solution by showing that and how the goal can be achieved; a dissolution by showing that the goal is not an important one to achieve; and a sceptical renunciation by showing that the goal cannot be achieved. With this structure in place, I am in a good position to highlight the relations between various positions in the literature.

There are broadly two camps on the issue of how Bayesianism and induction relate. The first, call them the *inferentialists*, take the goal of induction to be that of supplying a logic of inductive inference. According to them, Bayesianism solves the problem of induction in doing just that. The second, call them the *superbaby theorists*, take the goal of induction to be that of justifying an agent's credences about the unobserved at the time of analysis. They argue that this must be done in two steps: firstly, by determining the credences we ought to have had prior to having any evidence (that is, when we were superbabies[2]), and secondly, by determ-

[2] Hájek (ms) credits this term to David Lewis.

ining the way in which we ought to have reacted to the acquisition of evidence after that. According to them, Bayesianism makes progress on the problem of induction through its defence of conditionalisation, but not progress enough: the problem of the priors remains to be solved.

In this paper, I argue against both of these takes on the problem of induction (§1–2). Then, I tie together strings from both sections to put forward my own view, on which the goal of induction is that of justifying credences about the unobserved (though not specifically at the time of analysis), but on which a solution requires neither a determination of which priors one ought to have had, nor a determination of how one ought to update one's credences in light of evidence. I draw out some consequences of my account (§3), and I conclude with a brief summary (§4).

1.

The first group of Bayesians interested in induction are the *inferentialists*, according to whom the goal of induction is that of justifying a particular form of inference, called *inductive inference*. The most developed inferentialist position is that of Howson (2000), who presents the goal of induction as that of justifying "the process of reasoning that leads people to conclude that observational data obtained in suitably rigorous ways confirm some general hypothesis" (p. 6). In sum, inferentialists seek a sound logic of induction; that is, they seek to determine the correct inductive-logical relation between particular premises and conclusions.

There are several views on the nature of inductive inference. The first view holds that inductive inference is *enumerative* inference.[3] On this view, the premises of an inductive inference are propositions ascribing a single property to several objects of the same type, and the conclusion is that the next object of this type will also have the property in question. For example, the premises of such an inference might be propositions to the effect that ravens 1–99 are black, and its conclusion might be the proposition that raven 100 too is black. The second view

[3] For a prominent defender of this view, see what Jackson (1975) calls the *straight rule*.

of inductive inference holds that it consists in *generalising* inference.[4] On this view, the premises of an inductive inference are particular instances of a general proposition, and that general proposition is its conclusion. For instance, if the premises of such an argument are, again, that ravens 1–99 are black, its conclusion might be that all ravens are black. The third view holds that inductive inference is *ampliative* inference.[5] On this view, what characterises inductive inferences is that their conclusion is not logically entailed by their premises: the conclusion goes beyond what is contained in the premises. Note that both enumerative and generalising inferences are of this kind. In this paper, I shall adopt this third view of inductive inference, because alongside most Bayesians, I think that the problem is the same whether one reasons to the colour of the next raven, of the next seven ravens, or of all the ravens.

Inferentialists take Bayesianism to solve the problem of induction so construed. As Norton puts it: "proponents of Bayesian confirmation theory believe that they have the solution to a significant, recalcitrant problem in philosophy of science. It is the identification of the logic that governs evidence and its inductive bearing in science" (2011, p. 391). This argument is made most clearly by Howson (2000), and goes as follows. At the core of Bayesianism, there is a principle known as *conditionalisation*, which describes how agents ought to revise their credences in light of new evidence. Calling the credences that the agent has before receiving evidence her *pre-credences*, and the resulting credences her *post-credences*,[6] the principle dictates what post-credences an agent ought to have, given particular pre-credences and new evidence. According to Howson, this principle gives a logic of inductive inference—it specifies which conclusions agents ought to infer from

[4] This is the view expressed by Howson in his presentation of the problem of induction, cited above.

[5] This is the view endorsed for instance by Okasha (2001).

[6] What I have called pre-credences are usually known as *priors*, and what I have called post-credences are usually known as *posteriors*. But I have deviated from this common usage because the terminology of *priors* in the Bayesian literature is highly confusing: as we will see in §2, the term is used in a different way by superbaby theorists.

which premises—and as such, it constitutes a solution to the problem of induction. And furthermore, there exist numerous convincing defences of conditionalisation: among the most popular, we find decision theoretic arguments (Savage, 1954), Dutch Book arguments (Teller, 1973; Lewis, 1999), an expected accuracy argument (Greaves and Wallace, 2006), a conservativeness argument (Dietrich et al., 2016), and an accuracy dominance argument (Briggs and Pettigrew, 2020).[7]

Now, of course, Howson is aware that conditionalisation on its own, or even coupled with a fixed body of evidence, does not determine the conclusion that ought to be drawn from the inference: one must also specify pre-credences. But he denies that this is a problem, with a comparison to deductive inference. He cites Ramsey's claim that "this is simply bringing probability into line with ordinary formal logic, which does not criticise premises but merely declares that certain conclusions are the only ones consistent with them" (1926, p. 91). In other words, a sound inference takes one from a set of premises to the correct conclusion, but it does not determine a specific conclusion, because the premises are supplied exogenously. Just like we do not expect deductive inferences to yield a "categorical assertion" but a "conditional statement" (pp. 171–172)—so too we must expect a conditional statement and not a categorical assertion from a sound inductive inference. So on the inferentialist view, conditionalisation provides us with a logic of inductive inference, in which pre-credences figure as premises.[8]

Thus Howson concludes his book: "We have solved Hume's problem in about the only way it could be solved, by divorcing the justification for inductive reasoning from a justification of its consequences. Inductive reasoning is justified to the extent that it is sound, given appropriate premises." (p. 239, emphasis removed). He adds: "this is a book about *logic*, not rationality" (p. 239, emphasis original).[9]

[7] For an overview and critical discussion of some of these arguments, see Pettigrew (forthcomingb).

[8] The attempt to defuse the problem of induction by pointing to analogous problems with deduction, maybe along the lines of Carroll (1895), can be found in Haack (1976), and more recently, in Huber (2017).

[9] This is admittedly an odd conclusion for a book subtitled *Induction and the Justification*

But the pressing question concerns rationality, not logic. Johnson (1991) puts it very starkly:

> If we are to have beliefs about the unobserved, what ought we to believe? Should we expect the next emerald we examine to be green, or [bl]ue? [...] Why prefer one hypothesis about the unobserved to another, on the basis of knowledge merely of the observed? Why is the fact that the sun has always risen in the (observed) past a reason to expect it to rise tomorrow? Why is it not instead a reason to expect it not to rise tomorrow (or, indeed, a reason to expect whatever you please)? In a realm of thought where there are no *a priori* connections, what makes something a reason for believing one thing rather than another? (pp. 399-401)

So, contrary to the inferentialists' beliefs, the pressing goal is epistemic, not logical. This is not to say that inductive inference is irrelevant, but it is to say that it is subordinate: it is because we want to determine what the agent ought to believe, that might seek to study inductive inference. In fact, this dialectical move is already present in Hume. He says:

> If you were to ask a man, why he believes any matter of fact, which is absent; for instance, that his friend is in the country, or in France; he would give you a reason. [...] If we would satisfy ourselves, therefore, concerning the nature of that evidence, which assures us of matters of fact, we must enquire how we arrive at the knowledge of cause and effect. (1748, IV.4–5.)

Hume is clear: the goal to achieve—the goal of induction—is to determine "why he believes any matter of fact" about the unobserved, or again, to "assure us of [these] matters of fact". It is in order to determine these things that we must attend to inferences: the link between the "reason" for the belief and the belief itself, which Hume takes to be a matter of "cause and effect". So, what a closer reading of

of Belief.

Hume reveals is that the justification of inductive inference is subordinate to the justification of particular beliefs.

But this does not mean that the question of inductive inference is not central to Hume. Indeed, he argues that, as the matter of what to believe about the unobserved cannot be settled on the basis of observation, it must be justified on the basis of inference. He then considers two kinds of inference: deductive inference, and inductive inference. Deductive inference, he remarks, is too weak to take us from propositions about the observed to propositions about the unobserved. Inductive inference by contrast is not, but stands in need of justification. And any attempt at providing such a justification, he argues, will turn out to be unacceptably circular.[10] This is Hume's sceptical argument against the possibility of solving the problem of induction. Skeletally: one can only solve the problem of induction if one justifies inductive inference, and this cannot be done. Thus the following two stances are very close, but importantly distinct: Hume's stance that in order to achieve the (epistemic) goal of induction, one would have to provide a justification of inductive inference; and the inferentialists' stance that the goal of induction itself is to possess a justification of inductive inference. I suspect that the closeness of these two views is the reason for the inferentialists' mistake.

The closeness of the views opens up the possibility that, although inferentialists are mistaken in their identification of the goal of induction, they might nonetheless have made progress towards this goal. For if the task of determining the logic of induction is subordinate (in the sense described above) to the task of determining what to believe, and if inferentialists have succeeded in the former task, then they have thereby succeeded in the latter—that is, they have solved the (epistemic) problem of induction. The justification of inductive inference may not constitute the *goal* of induction, but according to Hume, it does constitute a (or even *the*) way of achieving that goal. This brings us to the second problem for the inferentialists. Hume writes: "The bread, which I formerly ate, nourished me; [...] but does it

[10] For a detailed presentation and critical discussion of Hume's arguments, see Stroud (1977) and Millican (1994).

follow that other bread must also nourish me at another time? [...] It must be acknowledged that there is here a consequence drawn by the mind; that there is a certain step taken; a process of thought, and an inference, which wants to be explained" (1748, IV.16.). Thus it is clear: for Hume, the inference which yields a claim about the unobserved is one that takes one from what one justifiably believes about the observed (that bread has so far nourished), to what one ought to believe about the unobserved (that bread will continue to nourish). But, as the inferentialists recognise themselves, and as I outlined above, the kind of inference that they justify does not share this structure. Indeed, Bayesian inference requires pre-credences as input, which, crucially, are about the unobserved—about, in fact, the very proposition under consideration. It follows that, unless one could provide a justification for these pre-credences, the resulting inference schema does not match Hume's, and inferentialists do not succeed at showing that the (proper) goal of induction can be met. This state of affairs has been recognised by many Bayesians, who grant that inferentialists are right in their claim that conditionalisation helps with induction, but insist that this is not sufficient: one must also specify which pre-credences are justified. I outline, analyse, and ultimately reject the picture endorsed by this second group of Bayesians in §2.

2.

I turn in this section to the second group of Bayesians interested in the problem of induction, whom I call the ***superbaby theorists***.[11] The superbaby theorists have correctly identified the goal of induction as epistemic. For example, they might wonder what credence an agent ought to have in the proposition that it will rain on her walk, given the evidence that she has, maybe that the weather has been quite bad these past few days, that it often rains in London, etc. Calling the credences an agent has at the time t_n of relevance to the theorist her ***nun-credences***, after the Greek νῦν for "now"; they think that the problem of induction is that of determining which nun-credences are justified.

[11] This view is made explicit for instance by Pettigrew (2016b) and Hájek and Lin (2017).

In order to determine which nun-credences an agent ought to have, the superbaby theorist reasons as follows. Given conditionalisation, the agent's credences at t_n are determined by her credences at t_{n-1}, together with the evidence E_n she acquired between t_{n-1} and t_n. Her credences at t_{n-1} are, in turn, determined by her credences at t_{n-2}, together with the evidence E_{n-1} acquired between t_{n-2} and t_{n-1}. Taking this line of reasoning to the limit, the agent's credences at t_n are determined by the evidence she has acquired by t_n, together with the credences she had at t_0, prior to having any evidence. To sum up: according to superbaby theorists, given a body of evidence at t_n, determining which nun-credences an agent ought to have—that is, solving the problem of induction—requires a justification of conditionalisation and an account of what credences the agent ought to have had prior to having any evidence (when she was, in Lewis' terms, a "superbaby"). These latter credences are typically called the agent's *priors*; and correspondingly, the problem of determining which such credences are justified is called the *problem of the priors*. However, as we saw in §1, the term "priors" is used in two different ways by Bayesians: sometimes to refer to the epistemic state against which a learning event happens, and sometimes to refer to the credences of a superbaby. To avoid this conflation, I will call the latter the agent's *proto-credences* for now.

Therefore, the superbaby theorist has provided an account of how the problem of induction would have to be solved: the rational proto-credences would have to be determined, and conditionalisation would have to be justified. It follows that the Bayesian defence of conditionalisation goes some way—but not all the way—to solving the problem of induction. Lange (2011) writes:

> Suppose it could be shown [...] that rationality obliges us to update our opinions by conditionalisation. [...] Such an argument would still be far from a justification of induction. [...] All depends on the prior probabilities plugged into Bayesian conditionalisation along with our observations. So in order to explain why we must reason inductively, [...] the rationality of Bayesian conditionalisation would have to be supplemented with some constraints on acceptable priors. (p. 80)

Even less optimistically, Okasha (2005) writes:

> Bayesian theory's constraints in no way tell us what our state-of-opinion should be, for any given set of evidence. Requiring an agent to update her probability function by conditionalising on new evidence only determines her posterior probability function given her prior probability function, and Bayesian theory is silent on what the latter should be, beyond demanding that it be coherent. [...] So the positive response to Hume [...] looks simply unavailable to the Bayesian. Hence widespread consensus among Bayesians that their theory leaves Hume's problem essentially untouched. (p. 184, emphasis original)

Let us illustrate this point with an example. Suppose that an agent considers whether bread nourishes at time t_1 and at time t_2. Let b_i represent the proposition that bread nourishes at time t_i. Now, assuming that the agent considers nothing but this matter, she considers four atomic propositions—four mutually inconsistent ways the world might be, given by a through d below. Let us assume that the agent learns by t_n that b_1 is true.

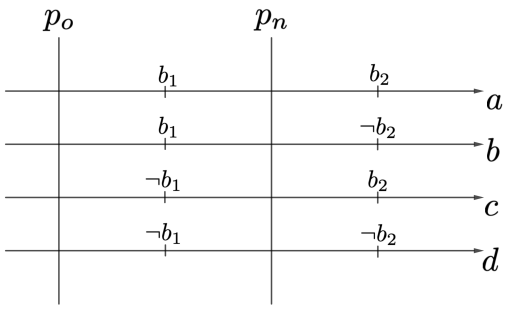

1. Illustration of the bread case

Suppose that this agent assigns equal proto-credences to a through d: she is such that $p_o(a) = p_o(b) = p_o(c) = p_o(d) = 1/4$. As this agent learns b_1, and given that she updates her credences by conditionalisation, her nun-credences ought to be such that $p_n(b_2) = p_n(\neg b_2) = 1/2$. But suppose that her proto-credences had

instead been such that $p'_o(a) = 1/3$, $p'_o(b) = 1/6$, $p'_o(c) = 1/3$, and $p'_o(d) = 1/6$. Then, given conditionalisation, her nun-credences ought to be such that $p'_n(b_2) = 2/3$ and $p'_n(\neg b_2) = 1/3$. The nun-credences it is rational for her to have depend on her proto-credal distribution. According to the superbaby theorist, in order to solve the problem of induction (in order to determine what nun-credences are rational), one must solve the problem of proto-credences (one must determine what credences superbabies ought to have).

In what follows, I will argue against the superbaby theorist's claim that, in order to solve the problem of the nun-credences, one must solve the problem of the proto-credences. In other words, I will argue that if one can solve the problem of induction, one can do so without appealing to proto-credences. Firstly, I will argue that, given conditionalisation, justifying the relevant credences at any time justifies them at every other time, such that there is nothing special about proto-credences. But secondly and more importantly, I will argue that justifying nun-credences is hard for exactly the same reason that justifying any other type of credences is hard. It follows that the strategy of derivatively justifying the agent's nun-credences by justifying some antecedently had credences is misguided, in the sense that it is completely unnecessary. It also follows that the problem of induction is broader than the problem of nun-credences.

Let us begin by highlighting what the problem is with nun-credences, using the toy example above. What is it, that is both difficult and important to determine? Suppose that, at the time t_n of analysis, the agent's evidence is that bread nourished at t_1. Then, her nun-credences ought to be such that $p_n(b_1) = 1$ and $p_n(\neg b_1) = 0$. Determining this much is in no way hard, so to the extent that there is a problem with nun-credences, it cannot be this. Neither is it is hard to determine what epistemic attitude to have towards c and d: since they have been ruled out by the evidence, they must each be assigned credence 0. No, what is hard is to determine whether bread will continue to nourish, or whether it will cease to nourish. That is, what is hard is to determine the relative confidence to have in a and b. And it is important to determine this; indeed, whether it is rational for an agent to

reach out for bread if she is hungry depends on whether she considers it likely that the bread will in fact nourish her. So, the problem of nun-credences, is that of determining the relative confidence to have in a and b. What do proto-credences and conditionalisation have to do with this?

According to the Bayesian who accepts conditionalisation, learning is a kind of deleting: it is the deleting of possibilities. The agent starts off considering a number of propositions, and as she gains evidence, she removes credal weight from those propositions her evidence entails are false. Then, she redistributes the weight among possibilities that have not been ruled out, in such a way that credal ratios are preserved. (This property of conditionalisation is called *rigidity*.) So, in our toy example, learning b_1 amounts to ruling out possibilities c and d, and redistributing credences across a and b in such a way that the confidence ratio in a and b remains the same. But now we can see why the superbaby-theoretic approach to the problem of nun-credences is misguided. As I explained above, what is hard about determining which nun-credences to have—what the problem is with nun-credences—is what relative confidence to have in a and b. But if learning takes place by conditionalisation, as the superbaby theorist assumes it does, the relative confidence an agent ought to have in a and b remains unchanged throughout her epistemic life, until one or both is ruled out by evidence. (Suppose that the agent went on to learn that a is true—the credence she would then ought to have would be such that $p(a) = 1$ and $p(b) = 0$.) This means that there is no normative primacy of the proto-credences: for the Bayesian who accepts conditionalisation, solving the problem of the relative credences to have in a and b at one time automatically solves it at all times. As such, the superbaby theorist is wrong to think that one must solve the problem of proto-credences in order to solve the problem of nun-credences.

But the wrongness of the superbaby theorist runs deeper still: it is not just that there is nothing special about proto-credences given conditionalisation, but that what makes it hard to determine what credences to have about the unobserved is the same at any time. Indeed, the problem of the nun-credences is that of

determining the relative confidence one ought to have in a and b—how confident one ought to be that bread will continue to nourish, as opposed to cease to nourish. It is a problem at the time of analysis because the agent's evidence (that which can be used to rule out possibilities) is silent on this matter: all her evidence is, by definition, about the observed. And it is a problem at the initial time for exactly the same reason: the question of which of a and b obtains is about the unobserved—as is every question for the superbaby—and so it not obvious how one should settle it. So, there is a sense in which the problem of nun-credences is identical to the problem of proto-credences: it is difficult to achieve the goals involved for precisely the same reason, namely, that doing so requires determining what to believe about possibilities about which one has no evidence. Here, we have arrived at the normative core of the issue. Whether the credences at hand are had at the beginning of an agent's epistemic life, or at the time at which she is being theorised, or at any other time, is besides the point. Rather, the point is that, for some propositions (the propositions about the observed), it is easy to determine what credences the agent ought to have in them: it is entailed by the agent's evidence. But for others propositions (the propositions about the unobserved), it is very difficult: the agent's evidence does not constrain the credences she ought to have in them at all. The superbaby theorist is right to think that the problem of induction is epistemic, but wrong to concern himself with the time at which the credences are had—the problem of induction is that of justifying credences about the unobserved, regardless of the time. Thus we might say, either that the problem of proto-credences and the problem of nun-credences are identical (in the sense that what makes them problems is identical), or we might say that they are both instances of a more general problem, namely the problem of induction (so the problem of proto-credences is the problem of induction applied to superbaby credences, and the problem of nun-credences is the problem of induction applied to current credences).

The superbaby theorist might object to this conclusion in the following way. The agent has no evidence when she is a superbaby, but she has plenty of evidence

at the time of analysis. So, what makes each of them problems is different: the problem of proto-credences is that of determining the credences she should have in the absence of evidence, and the problem of nun-credences is that of determining the credences she should have given her evidence at that time. But I think that this dichotomy between presence and absence of evidence is misleading. For, sure, the agent has evidence at t_n: she is certain that $a \vee b$ is true—but that in no way determines the relative confidence she ought to have in a and b individually. So, the problem of nun-credences is about what to believe given one's evidence, in the sense that, one can grant as much evidence as one wants about other propositions, and the problem of nun-credences still remains. This is because, by the very definition of the problem of nun-credences, the goal is to determine the credences an agent ought to have in propositions about which the evidence is silent. Although the agent may have evidence at t_n and not at t_o, she at no point (by stipulation) has evidence about which of a and b obtains. So, it is for the same reason that it is hard to determine the credences the agent should have at t_n and at t_o. The problems are identical; or alternatively, they are both timed instances of the problem of induction.

To conclude, the superbaby theorist is correct in his characterisation of the problem of induction as epistemic. However, his claim that the appeal to proto-credences is indispensible to the formulation of the problem of induction is false for two reasons. Firstly, given conditionalisation, determining what to believe about the relative plausibility of a and b at any time in the learning process would determine it for every other time until observation. It follows that proto-credences do not play a special role in addressing the problem of nun-credences. Secondly, and more importantly, determining what to believe about the relative plausibility of a and b is hard for exactly the same reason at all times—because any certainty one might have about the observed does not constrain credences about the unobserved. It follows that the problem with proto- and nun-credences is the same, namely, that of what credences one ought to have towards the unobserved. This, I submit, is the problem of induction properly understood.

3.

In my discussion of inferentialism, it transpired that the problem of induction is epistemic. In my discussion of superbaby theory, it transpired that the problem of induction is that of determining the credences that agents ought to have towards propositions about the unobserved. In this section, I draw out consequences of this account for four topics: briefly concerning conditionalisation, superbabies, and the priors; and at greater length concerning the subjectivism/objectivism debate in Bayesian epistemology.

1. Conditionalisation. As we saw in §1, inferentialists believe that the problem of induction is the problem of justifying conditionalisation. And as we saw in §2, superbaby theorists believe that the problem of induction is composed partly of the problem of justifying conditionalisation. So, although both camps disagree on many issues, they agree on this: to the extent that arguments for conditionalisation are successful, they constitute progress towards a solution of the problem of induction. (For inferentialists, complete progress; for superbaby theorists, merely partial progress.) But I have argued that the problem of induction is that of justifying credences about the unobserved, and that strategies appealing to diachronic considerations are misguided. It follows that conditionalisation is irrelevant to the problem of induction. This entails, *contra* both inferentialists and superbaby theorists, that Bayesians have made no progress towards addressing the problem of induction in their arguments for conditionalisation. Put another way, no matter their views on conditionalisation, all Bayesians are on a par when it comes to induction—even those who reject conditionalisation in favour of another updating rule,[12] and those who reject the existence of updating rules altogether.[13]

2. Superbabies. Superbabies do not exist—there are no agents who consider and have credences in a full algebra of propositions, but who have acquired no evidence. Superbaby theorists typically brush this concern off with a classic Bayesian move: the appeal to idealisation. Sure, there might not be any actual superba-

[12] Williamson (2010), Gallow (forthcoming).

[13] Hedden (2015)

bies, but this doesn't mean that we shouldn't model agents as if there were: it is a useful modelling assumption. The question of when idealisation is justified (if ever) is a complex one, which I cannot hope to settle here. I will note however that most Bayesians think that some but not all idealisations can be justified (for instance, no Bayesian accepts the idealisation assumption that agents are certain of all empirical truths). This suggests the need for a line to be drawn between acceptable and unacceptable idealisations, preferably underpinned by principled reasoning. In "Why Subjectivism?", for instance, I argued that (poric) Bayesians should welcome idealisations assumptions according to which agents have their means of inquiry perfectly. And on this account, the assumption that agents begin their epistemic lives as superbabies is not warranted. As far as I am aware, there does not exist a defence of a line between legitimate and illegitimate idealisations such that superbaby idealisations are legitimate. This constitutes a challenge to superbaby theorists: how are we to address the problem of induction if doing so requires a spurious appeal to superbabies? But if, as I have argued, the problem of induction is that of justifying credences in propositions about the unobserved, the appeal to superbabies is superfluous, and the challenge disappears. This is good news for Bayesians: their capacity to address the problem of induction does not rely on the contentious assumption that superbaby idealisation is legitimate.

3. The Priors. Bayesians commonly employ a particular strategy to push the problem of induction to the background. When considering the credences that an agent ought to have at the time of analysis, it is frequent to see Bayesians "assuming that the problem of the priors has been solved", or otherwise dodging issues by insisting that they depend on how one addresses the "problem of the priors", as though this were a distinct problem. But once we recognise that it is not—that in fact, the problem of determining the credences one ought to have prior to possessing any evidence is the same problem as that of determining the credences an agent ought to have at the time of analysis—this move is no longer warranted. It follows that the problem of induction is more pressing and less remote than Bayesians take it to be: it appears every time a theorist wonders what

credence an agent ought to have in a proposition about the unobserved. This, I hope, will prompt Bayesians to pay renewed attention to the problem of induction.

Furthermore, this consideration elicits an examination of the word "priors". I have insisted in §1 and §2 on the fact that the term is polysemous: it is used to refer both to the credences against which a learning event takes place (the pre-credences) and to the credences of an agent before the acquisition of any evidence (the proto-credences). But, as I have argued, mention of the "problem of the priors" (the problem of proto-credences) often serves as a way of unduly pushing the problem of induction back and away. Instead, the problem really is of all credences about the unobserved, and proto-credences are in no way special. So, I suggest that we should stop talking about them, and reserve the term "priors" for pre-credences instead.

4. Subjectivism/Objectivism/etc. If the problem of induction is the general (that is, arising at all times in the agent's epistemic life) problem of determining the credences an agent ought to have in propositions about the unobserved, it follows that the debate between subjectivists, objectivists, and other alternatives is not, as is commonly thought, a problem about the "priors", but instead, it is a problem that arises at any time an agent considers such propositions. Let us see why. The distinguishing mark of objectivism is the adherence to a principle of indifference: a principle according to which, in the absence of evidence, an agent's credences ought to be equally distributed among inconsistent propositions.[14] By contrast, the distinguishing mark of subjectivism is the adherence to the claim that, in the absence of evidence, no credence in a particular proposition is more justified than any other—all credences are equally justified.[15] It should be clear that these positions apply to proto-credences: since the agent as a superbaby has no evidence about any (non-trivial) proposition, objectivism and subjectivism make

[14] An indifference principle was first formulated by Bernoulli (1837/1954) and comes to Bayesians through Keynes (1921) and Carnap (1950). Prominent objective Bayesians today include Williamson (2010) and Pettigrew (2016a,b).

[15] As far as I'm aware, the only contemporary defence of subjectivism is in "Why Subjectivism?".

claims about the justification of particular credences in all such propositions. An objectivist and a subjectivist would disagree about which credences our toy agent of §2 ought to assign to propositions a through d: whereas the former would insist that she ought to assign credence $1/4$ to each, the latter would allow any credence function, so long as it is probabilistic. But the two positions can also be applied to our agent's nun-credences: the objectivist insisting in credence $1/2$ in both a and b; the subjectivist allowing any probabilistic credences in them. Indeed, we must reconceive of objectivism and subjectivism (as well as any position that purports to be an alternative) as answering the question: what credences ought agents have in propositions about the unobserved? Since this question can be asked irrespective of whether the agent entertains propositions about the observed, the possible answers might hold, not just of proto-credences, but of any credence function at any time.

So, objectivism and subjectivism are both ways of addressing the problem of induction. What kind of addresses are they? Objectivism is a solution to the problem, in the way outlined in the introduction: it is an answer to the question of how to achieve the difficult goal of assigning credences to propositions about the unobserved. It asserts that all propositions about the unobserved should be given equal credal weight. By contrast, subjectivism is a sceptical renunciation in the face of the problem, or what we might call an *inductive-sceptical* position: it asserts that there is no particular solution to the problem of induction—no credence in a proposition about the unobserved that is ever more justified than any other. This contradicts the widely held view that subjectivism "leads to" inductive scepticism. If one believes, as superbaby theorists do, that subjectivism is a position about which proto-credences are justified and that inductive scepticism is a claim about which nun-credences are justified, it makes sense to think of subjectivism as entailing inductive scepticism, given conditionalisation. But if what I have said is right, and subjectivism is a position about any credences about the unobserved, decoupled from conditionalisation, then it ceases to "lead" to scepticism, and instead becomes an inductive-sceptical position in itself. In turn, this entails that

arguments for subjectivism, such as the one I present in "Why Subjectivism?" are sceptical arguments—arguments for an inductive-sceptical position.

Given that I argued, in "Why Subjectivism?", for subjectivism, and given my positioning here of subjectivism as an inductive-sceptical position, I want to finish this section by discussing two widely held views about subjectivism seen through this lens. Firstly, some Bayesians claim that subjectivists/inductive sceptics find their position to be a happy state of affairs. For instance, Joyce (2011) writes that "according to the radical subjectivist, there is no problem of the priors. [...] Consistency only demands that priors obey the laws of probability: everything else is a matter of 'inductive taste'" (p. 18). Or again: "Coherent beliefs can be absurd, but subjective Bayesians tolerate them no matter how crazy they get, (short of violating the laws of probability). Many find this objectionable." (p. 30) But it is perfectly consistent to think that there exists no justification for any belief about the unobserved over any other, and that this is unfortunate or even "objectionable". Indeed, I cannot imagine a subjectivist who does not lament that subjectivism is true. This is in part because Bayesians of all stripes are committed to the fact that credences guide rational action. The fact that I am confident that bread will nourish me is what makes it rational for me to reach out for the bread when I'm hungry. So as Hume puts it, if inductive scepticism were true,

> [...] we should be entirely ignorant of every matter of fact, beyond what is immediately present to the memory and senses. We should never know how to adjust means to ends, or to employ our natural powers in the production of any effect. (1748, V.-6.)

Subjectivism does not entail that all credences should be "tolerate[d] no matter how crazy they get", it entails that no credence about the unobserved is more crazy than another, in the sense that it is impossible to settle whether a proposition is true, or even likely to be true, before its truth-value has been determined by observation.

Secondly, some Bayesians argue that, if subjectivism/inductive scepticism is true, there are specific epistemic attitudes agents should have. For instance, Okasha (2001, 2003) argues that "the sceptic demands that we suppress our background

beliefs, adopt the notorious 'theoretically barren' starting-point, and try to justify our response to the data without the help of those beliefs" (2001, p. 321). Lange (2002), in his response to Okasha, argues that "if the sceptic is asked to recommend a prior probability, he should suggest a distribution that makes no probability assignment at all to any claim about the world which concerns logically contingent matters of fact. (By this, [...] I mean making no assignment at all to any such claim, unless we have directly observed it to hold.)" (p. 228). But the subjectivist, *qua* inductive sceptic, is in no way committed to either of these claims. The claim to which she is committed is that no epistemic attitude is better than any other. This does not entail that agents should have no credences about the unobserved, or that they should have "barren" credences. It is consistent for instance with the claim that agents may adopt any credence function they please—a credence function that might be entirely un-barren, such as one according to which it is highly probable that bread will cease to be nourishing. That no credence function is justified does not entail that an agent cannot in fact have a particular credence function—it just implies that she would not be justified in having this credence function over any other. (This claim will turn out to be dialectically relevant for an argument I make in "The Nature of Awareness Growth".)

4.

By way of conclusion, let me sum up my arguments. I have been concerned, in this paper, with the relationship between Bayesianism and the problem of induction. There are two widespread views on this relationship, which I have called inferentialism and superbaby theory. I examined each of them in turn, and I argued that neither of them was tenable: the problem of induction is not that of determining the sound logic of inductive inference, nor is it that of determining what one ought to believe at the time of analysis specifically (doing which would require determining the rational proto-credences and justifying conditionalisation). Instead, I argued that the problem of induction is the general problem of what to believe about the unobserved. Once we accept this revisionary view, we are forced

to abandon the widespread belief that the defence of conditionalisation constitutes progress towards solving the problem of induction, we are faced with the fact that the problem is pervasive in Bayesian epistemology—and as such, more pressing than it is taken to be by contemporary Bayesians, and we are in a position to appreciate objectivism and subjectivism as ways of addressing the problem of induction—including the latter as an inductive-sceptical position.

3. The Nature of Awareness Growth

Abstract. Awareness growth—coming to consider propositions of which one was previously unaware—is a crucial aspect of epistemic thriving. And yet, orthodox Bayesianism cannot account for this phenomenon. As a remedy, two proposals have been put forward: I call them the *expansion* view and the *refinement* view. In the developing literature on this topic, there are no outright defendants of the refinement view: not only are there critics, but even those who endorse it do so half-heartedly. In this paper, I argue resolutely in favour of the refinement view. Furthermore, I show that its adoption greatly simplifies the answer to the question of how agents' credences should change upon awareness growth.

0.

Real agents often undergo what is known as *awareness growth*: they come to consider propositions of which they were previously unaware. Examples of this phenomenon abound: scientists formulate new theories, students discover new ways the world might be, people concede that they "hadn't thought of that". Importantly, such an epistemic shift does not seem irrational; on the contrary, it appears epistemically praiseworthy, and may even constitute a form of learning. This suggests that an adequate account of rationality must at least allow for awareness growth. But orthodox Bayesian epistemology does not: an agent's attitudes are defined on a single set of propositions, fixed throughout her epistemic life, and all rational

instances of learning take the form of credal change in response to evidence. This is the *problem of awareness growth*.[1] Given that Bayesians seek to give a complete account of rationality, this is a serious problem, for it threatens their project in a deep way: there is a seemingly rational phenomenon which is crucial to epistemic thriving that they rule out as irrational.

My aim in this paper is to show that there is a very good way to accommodate awareness growth in the Bayesian framework, or in other words, that there exists a good Bayesian model of agents undergoing awareness growth. As such, I seek to propose a solution to the problem of awareness growth. There are two possible options for modelling awareness growth: by *refinement*, or by *expansion* of the algebra. (I will give definitions of these options later on.) In the expanding literature on the problem of awareness growth, there are no outright proponents of the refinement view: not only are there critics, but even those who uphold it do so half-heartedly.

In this paper, I argue resolutely in favour of the refinement view. I begin by bringing the problem of awareness growth plainly into view, so as to state my thesis in full detail, and outline my assumptions. Then, I consider the case for the expansion view and against the refinement view, which I show relies on two claims, which I call the *contrivance* claim and the *defectiveness* claim (§1). I examine and refute these two claims in turn (§2–3). I conclude in favour of the refinement view, and I argue that this is good news as far as determining what to believe about newly considered propositions goes: it suggests an immediate answer to that question (§4).

[1] In the Bayesian philosophy of science literature, this problem was labelled by Glymour (1980) as the *problem of new theories*. Here, I use the term from formal epistemology because it is more general: it encompasses not only the formulation of new scientific hypotheses, but also more mundane phenomena which are identical in structure from the Bayesian perspective.

1.

In this section, I aim to get the problem of awareness growth and this paper's thesis fully into view. Thus the plan for this section is to describe what awareness is, what growth of awareness is, and what the problem is with awareness growth.

The first thing to note is that, in the Bayesian framework, an agent's epistemic state is represented by a probability triple $\langle \Omega, \mathcal{A}, p \rangle$, where Ω is a non-empty set, $\mathcal{A}$ is a Boolean algebra of Ω, and $p : \mathcal{A} \to [0, 1]$ is a probability function defined on $\mathcal{A}$.[2] The set $\mathcal{A}$ is interpreted as the set of propositions that the agent considers, that is, of which the agent is *aware*. As such, it will be our central concern in this paper. The function p assigns to each proposition in $\mathcal{A}$ a number between 0 and 1, which is interpreted as the agent's degree of belief, or *credence*, in that proposition.

In this paper, I will work with a particular interpretation of the agent's epistemic state: the ***map account***, named after Ramsey's famous assertion that belief is the "map by which we steer"[3] (1926). This is the dominant account of epistemic states, which is very nicely laid out by Yalcin (2018): an agent's epistemic state "represents the world as being a certain way (the map aspect), and [...] plays a certain role in the explanation of the agent's actions (the steering aspect)" (p. 24). Thus there are two parts to the account. Its representational part says that epistemic states represent the world as being in a particular way, and its dispositional part says that it is this representation of the world that explains dispositions to act. So on this account, the propositions of which the agent is aware (that is, the propositions included in the agent's algebra) are those about which the agent has an opinion which is linked to action. To illustrate this, consider the following example. An agent is about to receive a book of fiction. She considers that it might be a novel, or a book of poems, but has never encountered the genre of short stories. On the map

[2] On an alternative definition of probability functions, they are defined onto the reals, and not just $[0, 1]$.

[3] Although the phrase comes from Ramsey (1926), he would not have agreed with the account I spell out. Instead, he would have agreed with a purely dispositional account, such as the one endrosed by Mahtani, discussed below.

account, the agent is aware of the propositions *novel* and *poems*, but unaware of the proposition *stories*. That is, she has an opinion about whether the book is a novel, and an opinion about whether it is a book of poems, and she has dispositions to act on these opinions, but she does not have an opinion about whether it is a book of stories.

This account is importantly different from two accounts that can be found in the literature on awareness growth. The first is the *conscious awareness account*. Bradley (2017a) characterises propositions of which the agent is unaware as those which "are not available to the agent's consciousness at the time at which she is deliberating on it" (p. 253). One possible interpretation of this claim is that an agent is aware of a proposition just in case they are consciously entertaining that proposition at the time of choice. In line with this, one of the illustrations that Bradley gives of unawareness is the state of affairs whereby one "fails to recall [a particular prospect] at a particular time because it slips one's attention" (p. 253). But this cannot be the right interpretation of awareness. Consider a thirsty agent who picks up a glass of water while thinking about something else. By assumption, the agent does not consciously entertain the proposition that there is water in the glass. But it does not follow that he is unaware of it. Indeed, he may have a (stored) representation of the glass as containing water, even if his attention is not directed towards it, and this explains the rationality of his action. The agent does believe—and therefore is aware of—the proposition that there is water in the glass, even if he does not consciously entertain it at the time of choice.

Another interpretation of Bradley's claim is that an agent is aware of a proposition just in case they could consciously entertain that proposition, for some sense of *could*. On this spelling out of Bradley claim, we might count our water drinker as aware of the proposition that there is water in the glass, because he could direct his attention to that proposition; thus the proposition would be "available" to his consciousness, but in a looser sense. But I submit that facts about what an agent can consciously entertain are normatively irrelevant. It is widely accepted that non-biological entities, such as AI systems, may have epistemic states. Bradley

himself talks about thermostats as having beliefs (p. 63). This makes sense on the account of beliefs as maps by which we steer: the thermostat represents the world as being in a particular way (the room as having a particular temperature), and is disposed to act on this basis (is disposed to, e.g. turn on the heating if the temperature is below 19°C). But it does not make sense on the account of belief as requiring availability to conscious awareness: there is no sense in which a thermostat could consciously entertain a proposition. This illustrates a deeper point: what matters normatively is not contingent facts about an agent's cognition such as what thoughts do or might or could cross her mind; rather, it is the way in which she takes the world to be. If two people have identical representations of the world, they are identical as far as epistemology is concerned, even if they differ in some cognitive respects.

The second account of epistemic states I wish to contrast with the map account is the *dispositional account*, which Mahtani (forthcoming) discusses in relation to awareness growth. On this account, the dispositional aspect of the map account is preserved, but the representational aspect is discarded. Mahtani remarks that the account may be spelt out in at least two ways. On the first, the relationship between epistemic states and dispositions is one of constitution (to have a particular credence just is to have a particular disposition to act); on the other, it is one of causation (to have a particular credence is a causal factor in bringing about particular dispositions). She then shows that unawareness is impossible on this account, however spelt out. To get a sense of the shape of her argument, consider again the case of our book-lover. Presumably, this agent has a number of relevant dispositions. For instance, she would bet on the proposition that the book is a book of stories, were she presented with that bet. In exactly the same way, she would bet on all (mundane[4]) propositions, were she presented with a bet on them.[5]

[4] Mahtani restricts her claim to mundane cases of unawareness, and leaves aside cases such as transformative experiences (Paul, 2014) and childhood learning.

[5] Note that this does not assume that for each proposition, there is a unique amount of money that she would bet. More generally, it does not assume anything about what kind of dispositions the agent has. It merely assumes that she has dispositions.

It follows that she has epistemic attitudes towards *stories*, and in fact towards all propositions. There is no proposition of which she is unaware. Therefore, she cannot undergo awareness growth, and the associated problem does not arise.

Let us pause to reflect on the nature of the link between epistemic attitudes and dispositions. Mahtani assumes that this link is what we might call *positive*: credences and dispositions are related by nature. But this link could also be *normative*: credences and dispositions might be related by considerations of rationality. Whereas credences are or cause dispositions to act on a positive reading, they rationalise dispositions on a normative reading. To illustrate the difference, suppose that an agent does not want to get wet, and has a high credence that it is raining outside. On the positive reading, it would be metaphysically impossible for the agent to (be disposed to) go outside nonetheless. For to have a high credence that it is wet just is to not be disposed to go outside given that one hates getting wet. On the normative reading by contrast, it is perfectly possible for the agent to go out—but it is irrational. This agent has the wrong disposition to act, given his credences and desires. How is this distinction between positive and normative readings relevant to awareness growth? The best way to spell out the dispositional aspect of the map account is in a normative way. Indeed, on the map account, it is because of the agent's representation of the world that she ought to act in a particular way. Her high credence to the effect that it's raining makes a disposition to stay inside rational.[6] By contrast, the dispositional account of credence calls for a positive reading. The agent's dispositions are metaphysically cemented: for her to have a particular disposition simply entails that she has a particular epistemic attitude.

So, the normative reading allows for a metaphysical gap between epistemic at-

[6] It is hard to imagine a positive reading of the dispositional aspect of the map account. For it is presumably coherent for an agent to (a) not want to get wet, (b) have a representation of the world on which it is likely raining, and (c) be disposed to go outside. But if this agent is coherent, what does she believe? The representational aspect indicates that she believes it is raining, and the dispositional aspect (on a positive reading) indicates that she believes that it is not raining.

titudes and dispositions, which is absent on the positive reading. And, as Mahtani shows, it is hard to even make sense of the phenomenon of unawareness without this gap. If there is no such gap, then the dispositional account is plausible and awareness growth impossible. If we do allow such a gap, then a good account of epistemic attitudes is the map account, on which awareness growth is possible.[7] This tracks an old debate between the behaviourists, operationalists, and functionalists on the one hand; and the mentalists, judgmentalists, and representationalists on the other. (The fact that many Bayesians, starting with Ramsey, have been of the former camp helps to explain why, as a sociological matter of fact, little attention has been given to the problem of awareness growth.) In this paper, I will side with the latter camp, and assume that there is a normative link between credences and dispositions. This is in part because I believe it is correct (though I will not defend this here), and in part because it is a highly popular position held by many, if not most Bayesians today.

Now that we are equipped with an account of awareness, we can turn to the phenomenon of awareness growth. It will be useful to distinguish three types of awareness growth as thought of pre-theoretically. Consider again the agent who knows that she will receive a book of fiction as a present, and who considers that it might be a novel, or a collection of poems. There are three types of awareness growth that she might undergo. Firstly, she might come to consider that the book might be a paperback or a hardback. She would then consider four possibilities: paperback novel, hardback novel, paperback poetry, and hardback poetry. I call this *orthogonal* awareness growth: the newly considered propositions are logically independent from the old ones. Secondly, she might come to consider that, if the book is a book of poems, it might be a book of prose poems or a book of poems in verse. She would then consider three possibilities: novel, prose poems, and verse poems. I call this *internal* awareness growth: the newly considered propositions jointly entail a single old one. Thirdly and finally, the agent might come to consider

[7] Other accounts of credence with a dispositional aspect normatively read are possible in principle.

that the book also might be a collection of short stories. She would then consider three possibilities: novel, poems, and stories. I call this *lateral* awareness growth: the newly considered proposition is inconsistent with the old ones.

How can the Bayesian model these epistemic events? As mentioned above, in Bayesianism, the propositions that an agent has epistemic attitudes towards are represented by a Boolean algebra $\mathcal{A}$. There are two ways, in principle, to enlarge this algebra to another $\mathcal{A}'$. The first is by *refinement*: $\mathcal{A}'$ is a fine-graining of $\mathcal{A}$. The second is by *expansion*: $\mathcal{A}'$ contains all the atomic propositions of $\mathcal{A}$, and more. These two operations are illustrated in the figure below.

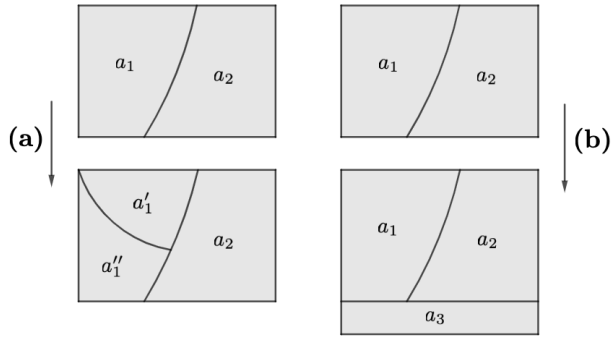

1. Refinement (a) and expansion (b) of an algebra of propositions

There is a consensus in the literature that orthogonal and internal growths of awareness should be modelled by refinement. So, for instance, the agent who comes to consider that the poems may be in verse or prose should be modelled such that the element *poems* of her old algebra $\mathcal{A}$ corresponds to the union of the elements *verse* and *prose* of her new algebra $\mathcal{A}'$. However, there is a dispute about how to model lateral growths of awareness. Should it be modelled by refinement (the *refinement view*) or by expansion (the *expansion view*)?

An early and influential discussion of awareness growth can be found in the work of Earman (1992).[8] Earman assumes that awareness growth would have to be modelled by refinement; that is, he does not even entertain the expansion view. He then recognises that the refinement view mandates the inclusion of what Shi-

[8] Relevant earlier work includes Glymour (1980) and Salmon (1990).

mony (1970) had called a *catch-all* proposition in the agent's algebra. Roughly, a catch-all proposition is a proposition which expresses something like "a possibility beyond those I have actively considered"; more precisely, it is the complement of the union of all "actively considered" propositions.[9] Why does the refinement view mandate the inclusion of such a proposition? Well, on the refinement view, awareness growth consists in the fine-graining of a previously entertained proposition. But the kind of awareness growth we are considering, lateral awareness growth, is such that the newly considered proposition is inconsistent with the old actively considered ones. So, there must have been an old proposition besides the actively considered one to be fine-grained: the catch-all proposition. Having recognised that the refinement view requires the inclusion of a catch-all proposition, Earman raises a concern about such an inclusion, which he takes to be fatal, and he concludes that Bayesians cannot model awareness growth. It is only two decades later, in the formal epistemology literature, that we find another take on the problem of awareness growth. Wenmackers and Romeijn (2016) continue to assume that awareness growth must happen by refinement, but do not take the aforementioned concern to be fatal to the view. Instead, they put forward a positive proposal to address this concern. By contrast, Bradley (2017a) and Steele and Stefánsson (ms) do take the concern about the refinement view to be fatal, but instead propose and argue for the expansion view, and as such hold that the problem of awareness growth can be solved.

We see that there is significant disagreement between authors in the literature. However, everyone agrees on one point: there is an important concern with the refinement view. What is it? The problem is that, unlike other propositions, it is uniquely difficult to assign a rational prior to the catch-all proposition. Earman (1992) claims that assigning such a prior would have to be done in a way that is "arational" (p. 197). In a similar vein, Bradley asks rhetorically: "given that we

[9] I have put "actively considered" in scare quotes because, although the distinction between actively considered propositions and catch-all propositions might be intuitively appealing, I will in fact argue in §4 that there is no normatively relevant such distinction. Nonetheless it is helpful to speak in this way for ease of exposition.

don't know anything about the prospects that we are potentially unaware of, on what basis are we to determine [...] what probability we should assign to the catch-all prospect?" (p. 255). Steele and Stefánsson (ms) write: "it is hard to see how an agent could assign a probability to a catch-all—the content of which she is assumed not to know—since it would be impossible for her to know whether her evidence speaks in its favour". Wenmackers and Romeijn (2016) write that "unlike the other hypotheses, it is not produced by a scientific theory, but rather it results from a meta-theory" (p. 1333), and they conclude that "since the catch-all is not based on a scientific theory, the usual 'arational' considerations [...] for assigning it a prior, namely by comparing it to hypotheses produced by other theories, do not come into play here" (p. 1234). According to Henderson et al. (2010), the refinement view "is an unsatisfactory solution since there is no particularly principled way to decide how much initial probability should be assigned to the catchall" (p. 190). The consensus is then that the catch-all proposition is defective: it is uniquely difficult to know what credence to assign to it. I will call this the *defectiveness claim*.

And besides a belief in the defectiveness claim, we can discern another belief, held in common between all these authors.[10] They hold not only that the inclusion of a catch-all is mandated by the refinement view, but the converse, too: that the refinement view is the only reason we would include a catch-all proposition in the agent's algebra. This is implied by the argumentative strategies of Shimony, Bradley, and Steele and Stefánsson: they take the presumed defectiveness of the catch-all hypothesis to be a reason to abandon that the refinement view, but no more. It is also implied by the shape of Wenmackers and Romeijn's dialectic: it is in order to maintain the refinement view that they address the defectiveness concern. So, not only are catch-all propositions defective, but their inclusion in an agent's algebra is contrived to uphold the refinement view. I shall call this second claim the *contrivance claim*.

[10] Besides the authors cited above, there are numerous other discussions of awareness growth and catch-all propositions. Examples include Hill (2010), Henderson et al. (2010), and Carr (2015). But none of these authors hold a view that departs from those outlined above, on the topic that concerns us.

In this paper, I seek to refute these two claims. Firstly, I argue that the inclusion of a catch-all hypothesis is mandatory for reasons independent of awareness growth, such that a proponent of the expansion view, and a proponent of the irrationality of awareness growth, would also have to include a catch-all proposition in agents' algebras. So, the contrivance claim is false (§2). Secondly, I argue that there is no difference between the catch-all proposition and the other propositions that would entail the impossibility of assigning a prior to the former and not to the latter. So, the defectiveness claim is false (§3). Finally, I conclude in favour of the refinement view, and, I consider a question related to the expansion/refinement debate: that of how one ought to change one's credences upon awareness growth. I show that the adoption of the refinement view simplifies the debate around this question (§4).

Before I embark on this project, let me lay out some assumptions. Firstly, I will assume that epistemology is a means-ends endeavour in the sense I presented in "Why Subjectivism?": the ultimate epistemic aim of an agent is to determine what is the case, and the agent inquires to achieve that aim. This fits nicely with the map account of credences. The agent seeks to match her credences with the world (the representational element), so as to make the best possible decisions (the dispositional element). In "Why Subjectivism?", I discussed two accounts of justification on a means-ends approach to epistemology: the *poric* account, on which an epistemic state is justified just in case it is warranted by the agent's means of inquiry; and the *telic* account, on which an epistemic state is justified just in case it relates in the right way to the aim of inquiry. In this paper, I will remain neutral between these two accounts, though it will be helpful to have the distinction in mind when mapping possible arguments for various positions. Secondly, I will assume that the trivial/non-trivial distinction between propositions (defined below for the unfamiliar reader) tracks, not the distinction between the logically necessary and the logically contingent, nor the distinction between the metaphysically necessary and the metaphysically contingent, but the (epistemic) distinction between the *a priori* and the *a posteriori*. So, $\emptyset$ denotes false propositions

that are determinable *a priori*, Ω denotes true propositions that are determinable *a priori*, and other propositions A_i denote propositions determinable *a posteriori*. Finally, I will make an idealisation assumption, to the effect that, when confronted with a proposition, (i) the agent is always in a position to determine whether it is determinable *a priori* or *a posteriori*, (ii) if the proposition is determinable *a priori*, the agent is always in a position to be certain of its truth-value, and (iii) if the proposition is determinable *a posteriori*, and it constitutes evidence for the agent, the agent is in a position to be certain of its truth-value. Such an assumption is common in Bayesianism, and is made throughout the literature on awareness growth.

2.

In this section, I argue against the contrivance claim. My strategy is as follows. I begin by considering a widely-discussed putative norm of rationality: *regularity*. I consider the two significant objections to regularity, and I argue that, whilst one can be set aside, the other should prompt us to abandon regularity in favour of a weaker norm: *humility*. I discuss a debate in the literature which pertains to this norm, and I propose a positive argument for it on that basis. I then show that humility entails that a catch-all proposition must be included in an agent's algebra, irrespective of whether she ever undergoes awareness growth. It follows that the contrivance claim is false.

Let us begin with regularity. Formally, regularity says that agents must ascribe non-extremal credences to non-trivial propositions. To see what this means, remember the Bayesian assumption that an agent's credences form a Boolean algebra $\mathcal{A}$. There are two kinds of propositions in a Boolean algebra: the so-called *trivial* ones, denoted Ω and $\emptyset$, and the *non-trivial* ones (all the others), usually denoted with letters. Regularity demands that the non-trivial propositions be assigned non-extremal credences; that is, credences that are neither 0 nor 1.

> *Regularity.* An agent's credence function p must be such that, for all $A_i \in \mathcal{A}$ such that $A_i \neq \Omega$ and $A_i \neq \emptyset$, $p(A_i) \neq 1$ and $p(A_i) \neq 0$.

So, regularity is the converse or complement of another Bayesian norm, trivial omniscience, which mandates the adoption of extremal credences in trivial propositions.[11]

> *Trivial omniscience.* An agent's credence function p must such that $p(\Omega) = 1$ and $p(\emptyset) = 0$.

Together, regularity and trivial omniscience entrench the distinction between trivial and non-trivial propositions. Not only are they different in kind, as is entailed by the requirements that the propositions an agent considers form a Boolean algebra; but they are different in that different kinds of credences are allowed in them—non-extremal credences in non-trivial propositions; extremal credences in trivial propositions.

There are two problems with regularity. The first and most widely discussed is that, in Lewis' words, "there are too many alternative possible worlds to permit regularity" (1980, p. 267). This objection is made famously vivid by Bernstein and Wattenberg (1969) and popularised by Hájek (2003). They instruct their reader to consider a dart being thrown at a dart board, and ask: what credence ought one assign to the proposition that the dart will land exactly in the centre of the board? What the rhetorical question shows is that, since there are uncountably many points at which the dart could land, if the agent wants to assign equal credence to all possible points at which the dart could land, she will have to violate probabilism. For there is no probabilistic function that can assign equal credence to uncountably many mutually inconsistent propositions. More precisely, the example shows that the following three claims are mutually inconsistent: that $\mathcal{A}$ may be uncountable, that credences must be real-valued, and that credences must be uniform over an algebra in the absence of evidence. As a matter of fact, I reject all three claims.[12]

[11] Probabilism is the conjunction of trivial omniscience and countable additivity, provided that credences are defined onto the reals.

[12] My argument against a prescription of uniformity can be found in "Why Subjectivism?"; my reason for rejecting the insistence that credences be real numbers is that I see no problem with infinitesimal probabilities—for an excellent overview, see Wenmackers

But I will not defend this here.

Instead, I will merely restrict our domain of concern by stipulation, to cases in which an agent considers finitely many propositions. The main reason for this is that the cases of interest as far as awareness growth is concerned are all finite. All real agents who wonder what genre a book of fiction will turn out to be consider finitely many options. All real agents who wonder which theory of fundamental physics is the correct one consider finitely many options. And so on. And given that the concern voiced by Lewis and Hájek does not apply to finite cases—for in those cases there is enough probabilistic credal mass to distribute among competing options—we can circumvent the worry. A second, more tentative reason, rests on the thought norms of rationality can be divided into two kinds: the substantive norms, which dictate the credences agents ought to have in particular propositions, and the formal norms, which dictate how agents' credences in various propositions fit with one another. Regularity is a substantive norm: take a non-trivial proposition—it tells you which credences you may or may not have in it. The argument that Lewis raises against regularity, by contrast, is formal: it concerns the ways in which our attitudes across propositions may or may not fit together. And, it seems to me inappropriate to reject a substantive norm on the basis of a formal problem. This consideration is what, I presume, moves so many Bayesians to admitting infinitesimal credences—see Wenmackers (2019) for an overview.

There is, however, a second objection to regularity in the Bayesian literature. The following norm is widely held by Bayesians:[13]

> *Evidential omniscience.* An agent's credence function p must be such that $p(E_i) = 1$ and $p(\neg E_i) = 0$, for all evidential propositions E_i.

(2019); and my reason for rejecting the possibility of an uncountable algebra is that all propositions in an algebra must be learnable by observation (they must be able to come to constitute evidence), and that, to use the dart board example, we cannot discriminate between uncountably many positions on the board.

[13] Conditionalisation is the conjunction of evidential omniscience and rigidity.

But, as Hájek (2012) points out, accepting regularity entails committing oneself, either to the rejection of evidential omniscience, or to the claim that there are no evidential propositions. Few people have been bothered by this, presumably because most reject evidential omniscience. However, there is a dialectical difficulty with this shrug. As I outline in "Why Subjectivism?", the reason for which many reject evidential omniscience is that they think it is too demanding: it requires of agents to be perfect observers—to always recognise of a proposition whether it is evidential. For instance, an agent may not be able to tell whether the proposition expressing that she has hands is an evidential proposition for her. But suppose that they could. Suppose that we conceded to agents more inquiring capacities than they in fact had, in such a way that they were always in a position to become certain of what they observed. Then, would we not want them, against regularity, to satisfy evidential omniscience? Now we are in a position to see the problem with the shrug. The reason we should want agents to satisfy regularity is not because they are deficient inquirers. Rather, whatever norm of rationality we end up endorsing, we should endorse because it pertains to proficient agents. And we have assumed, as part of a common idealisation assumption, that agents are perfect observers, as well as perfect reasoners. For these agents, adopting credence 1 in a proposition they have received as evidence is perfectly rational.[14] So, we need to weaken regularity to allow for this fact. Where $\mathcal{E} \subseteq \mathcal{A}$ is the set of evidential propositions, we get:

> ***Humility.*** An agent's credence function p must be such that, for all $A_i \in \mathcal{A}$ such that $A_i \neq \Omega$, $A_i \neq \emptyset$, and $A_i \notin \mathcal{E}$, $p(A_i) \neq 1$ and $p(A_i) \neq 0$.

Given the interpretation we have assumed of the trivial/non-trivial distinction, humility says that agents ought not to be certain of the truth-value of empirical propositions about the unobserved.

[14] An anonymous referee for FEW suggests considering regularity as a norm for prior credences only. But, as I have argued in "Bayesianism and the Problem of Induction", I think that this type of strategy is misguided.

This norm is appealing to many Bayesians, but it is surprisingly difficult to provide a convincing argument in its favour. Let us begin by attempting a poric argument. On a poric view of epistemic justification, an epistemic attitude is justified just in case it is warranted by the agent's means of inquiry. This is how I argued, in "Why Subjectivism?", for trivial and evidential omniscience. I made two assumptions: the dualism assumption, according to which agents have two means of inquiry—reason, and observation; and the idealisation assumption I also made above, according to which they have these means of inquiry perfectly. On this basis, I argued that Bayesian agents are in a position to determine the truth-value of *a priori* propositions, and of *a posteriori* propositions about the observed (evidential propositions), and so, ought to have the pertinent extremal credences in these propositions. At first sight, this might seem to provide a promising way of arguing for humility. For how could agents with dual means, even perfectly had, possibly be in a position to determine the truth-value of propositions about the unobserved?[15] It follows that it is not the case that agents ought to have extremal credences in *a posteriori* propositions about the unobserved.

But there is a distinction between it not being the case that agents ought to have such credences, and it being the case that agents ought not to have such credences. The poric argument above establishes the former, but not the latter claim; and humility is the latter. Might we be able to augment the poric argument in such a way as to establish humility? The only non-*ad hoc* way to do so would be to supplement the norm according to which they ought to have the credences warranted by their means, with a complementary norm according to which they ought not have any credences which are not warranted by their means. But this would

[15] One might protest that agents have another means of inquiry, beyond reason and observation: testimony. But we can make the argument work anyway—either by insisting that testimony is a sub-category of observation (I learn about what the experts believe through my senses), or by weakening humility further to a norm we might call *super-humility*, according to which agents ought to have non-extremal credences in all unobserved, untestified propositions—the proposition that the book will either be a novel or a poetry collection is of this kind.

entail that *all* credences about the unobserved are impermissible; that agents must assign no credence to any proposition besides *a priori* propositions and propositions about the unobserved.[16] But, setting aside the question of whether this is desirable or coherent, it breaks with the spirit of humility. For I take it that the driving intuition behind humility is that there is something special about extremal credences; that assigning credence 1 to *novel* $\vee$ *poems* is worse, in some sense, than assigning credence .9 to it. And if we adopt the complementary norm, we will not be able to mark this difference. So, I think a poric argument fails.

This brings us to the classic argument for humility, which is telic. It is a semi-Dutch Book argument, initially due to Shimony (1955). This argument goes like this: an agent with non-humble credences is liable to a loss (if the empirical proposition in which she is certain turns out false) but does not stand to gain anything (if it turns out true), and this situation is bad for her—she ought to avoid it. As Skyrms puts it, "there is a bet which we will consider fair even though we can possibly lose it but cannot possibly win it" (1980, p. 74). This can be made more vivid, by considering the fact that it would be rational for an agent with credence 1 in, say, the proposition that the book will be a novel or a collection of poems, to agree to getting nothing if the book is indeed a novel or poems, but to being tortured to death if it is not. But it is hard to say precisely what is wrong with agreeing to such a bet, in a way that is not circular. We cannot insist that the agent should not agree to it because she should not be certain: that is what we are trying to prove. And we cannot say that she shouldn't agree because such a bet would look bad from her own perspective: for by assumption, she is certain that the book is either a novel or a collection of poems. As Hájek says, "an omniscient God who knows which world is actual, and so concentrates all credence on that world, is semi-Dutch bookable, and none the worse for that!" (2012, p. 420).[17] The mere possibility of losing money cannot be sufficient to ground the claim of irrationality.

Rather, the problem with violating humility is that it stands in tension with

[16] I discussed this issue in "Bayesianism and the Problem of Induction".

[17] For further discussion of the semi-Dutch Book for regularity, see Easwaran (2014).

the very nature of epistemic inquiry. Let me explain. As I announced earlier, I am assuming a means-ends conception of epistemology in this paper, on which the agent deploys her means to achieve the ultimate epistemic aim, that of determining what is the case. Firstly, consider an agent who assigns credence 1 to a proposition A_x that is determinable *a priori*—for instance, the proposition that 2+2=4. Given our dualism and idealisation assumptions, our agent's means of reasoning allow her to determine that A_x is the case. Again: our agent has deployed her means in order to achieve an epistemic aim, namely, to settle whether A_x. Now, consider an agent who assigns credence 1 to a proposition A_y about the unobserved. This agent has also settled whether A_y: assigning credence 1 to it entails ruling out the possibility that it may be false. But not only has she not settled it on the basis of her means (for by assumption, her means of inquiry do not allow her to determine the truth-value of propositions about the unobserved), but she is in a position to know that (for again by assumption, she is able to determine of any proposition whether it is determinable *a priori*, determinable *a posteriori* and observed, or determinable *a posteriori* and unobserved). So, this second agent's mistake lies in having broken the normative link between the means and ends of inquiry: she has got to the end of inquiry without achieving it, one might say, since she has not gotten there through the exercising of her means. This stands in contrast with yet a third agent, who also considers A_y, and who assigns it credence .999. Although this agent has an extremely similar representation of the world, and ought to act in an extremely similar way to the previous agent, he has not settled whether A_y, and still takes $\neg A_y$ to be a live possibility. Thus, his epistemic state is perfectly consistent with a means-ends epistemology: he has not got to the end of inquiry. So, understanding extremal credences as the goal of epistemic inquiry allows to see why humility is true.

How is the truth of humility relevant to the contrivance claim? Remember, according to the contrivance claim, the only reason to include a catch-all proposition in an agent's algebra is to allow for awareness growth on the refinement view. But if humility is true, a less-than-fully aware agent must be modelled as having a catch-

all proposition in her algebra regardless of whether she ever undergoes awareness growth. Indeed, humility entails that an agent ought to assign a non-extremal credence p to $\textit{novel} \vee \textit{poems}$ (in virtue of the fact that it is an empirical proposition about the unobserved); and probabilism dictates that, in order to be rational, she ought to have a credence of $1 - p$ in $\Omega \backslash \{\textit{novel} \vee \textit{poems}\}$. But $\Omega \backslash \{\textit{novel} \vee \textit{poems}\}$ is, of course, a catch-all proposition: its content can be articulated as "the book I will receive is a book of fiction, but neither a novel nor a poetry collection". Thus all less-than-fully aware agents must be modelled as having a credence in a catch-all proposition, and the contrivance claim is false.

Before closing this section, I should make a remark about the views of two vocal proponents of the contrivance claim, Bradley (2017a) and Steele and Stefánsson (ms), who both reject the inclusion of catch-all propositions. I went to great lengths in this section to argue that humility is a norm of rationality, so I could use it against the contrivance claim. But I should note that neither Bradley nor Steele and Stefánsson believe that agents ought to be absolutely certain that a proposition like $\textit{novel} \vee \textit{poems}$ is true. Indeed, Bradley writes: "what should [the agent] do if she is conscious of the possibility *(as she should be)* that she may not be aware of all relevant prospects? (p. 254, emphasis mine). Steele and Stefánsson appeal, throughout their paper, to an *awareness context*, against which propositions may or may not count as "tautologies". So for instance, they consider an agent whose initial awareness context is such that $\textit{thriller} \vee \textit{comedy}$ is a "tautology", and whose awareness context changes so that $\textit{thriller} \vee \textit{comedy} \vee \textit{drama}$ becomes a "tautology". (The agent is wondering about the genre of the film showing at their local cinema.) Given trivial omniscience, the agent goes from rationally ascribing credence 1 to $\textit{thriller} \vee \textit{comedy}$, to rationally ascribing credence 1 to $\textit{thriller} \vee \textit{comedy} \vee \textit{drama}$. So, neither Bradley nor Steele and Stefánsson's rejection of catch-all propositions (nor their concurrent adherence to the contrivance claim) is motivated by the thought that agents may be absolutely certain of the truth-value of propositions about the unobserved.

Instead, they agree with the basic thought behind humility, that complete cer-

tainty in a proposition like *novel* $\vee$ *poems* is impermissible. But they handle this, not by mandating the inclusion of a catch-all proposition, but by divorcing absolute certainty from credence 1. Instead, credence 1 for them becomes something akin to "certainty in a context". In part because they do not give an account of awareness contexts, I cannot hope to settle whether this is a good interpretation of the Bayesian framework in this paper.[18] I will remark however that this move is inconsistent with both aspects of the map account of credences that I have been assuming. Firstly, it is inconsistent with the representational aspect of the account, according to which an agent's credences reflect how she takes the world to be. Of course, the possibilities that she is consciously entertaining, or the possibilities that are salient to her, can change according to context. The possibility that I could die in a plane crash is much more vivid to me when I am on a plane, than when I am at home with no plane journeys on the horizon. But whether I take it to be possible that I might die in a plane crash should not change according to context. The way one takes the world to be—one's map—must remain constant across contexts.[19] Secondly, the context-sensitivity move is also inconsistent with the dispositional aspect of the account, which maintains that an agent's credences rationalise her dispositions to act. Presumably, the authors who propose that move do not believe that it would be rational for an agent to bet her life on *novel* $\vee$ *poems*, given that, in Bradley's words quoted above, "she is conscious of the possibility (as she should be) that she may not be aware of all relevant prospects". But in order to vindicate this, they would have to sever the link between credence and rational action; that

[18] Though, see Mahtani (forthcoming) for an argument to the effect that their claims are internally inconsistent.

[19] Although they do not say this explicitly, Steele and Stefánsson's examples suggest that endorse the same account of awareness as Bradley does: the *conscious awareness* account, discussed and rejected in §1. On such an account, I might be aware of the possibility of a plane crash whilst on a plane, and unaware of it at other times. The problem with this, illustrated both above and in §1, is that there is a discrepancy between the way an agent takes the world to be, and the propositions that she is consciously entertaining, and that it is to the former that epistemology pertains.

is, to reject the dispositional aspect of the map account.[20] Given the importance for many Bayesians of the link between credence and rational action, this, I think, is an unappealing strategy.

Let me conclude this section with a brief summary. Given a means-aims conception of epistemology and the map account of credences, humility is a (true) norm of rationality, and so, the contrivance claim is false. Furthermore, those who embrace the contrivance claim and who reject the inclusion of catch-all propositions must reject the map account of credence.

3.

In this section, I argue against the deficiency claim in two stages. In the first stage, I analyse the argument given in favour of the defectiveness claim, and I show that it fails. In the second stage, I propose an argument against the claim. I then conclude that it must be abandoned.

According to virtually everyone who has written on awareness growth, including on the formulation of new scientific theories, there is a serious problem with catch-all propositions: it is uniquely difficult to know what credence to assign to them (Shimony, 1970; Glymour, 1980; Salmon, 1990; Earman, 1992; Henderson et al., 2010; Wenmackers and Romeijn, 2016; Bradley, 2017a; Steele and Stefánsson, ms). This is the deficiency claim. The argument that all these writers take to establish this claim is powerfully articulated by Bradley: "given that we don't know anything about the prospects that we are potentially unaware of, on what basis are we to determine [...] what probability we should assign to the catch-all prospect?" (p. 255). That is, what supposedly makes catch-all propositions uniquely difficult to assign justified credences to is that, unlike other propositions that the agent

[20] In an earlier chapter of his book, Bradley endorses what he calls the *rationality hypothesis*: "Rationality requires of an agent that she prefer one prospect over another if and only if the expectation of benefit conditional on the truth of the former is greater than the expectation of benefit conditional on the truth of the latter, relative to her degrees of belief and desire" (p. 21). This seems to me to assume the dispositional aspect of the map account of credence.

considers, we do not know what they say. Our book-lover knows what *poems* says, but she does not know what *fiction that isn't novel or poems* says. So, how could she assign a particular credence to the latter? Thus the common argument for defectiveness is built on two claims: (1) there is a principled difference between "considered" propositions and the catch-all proposition, and (2) this difference entails a difference in tractability. In what follows, I argue against both of these claims.

Let us begin with the claim, (1) that there is a sharp distinction between considered propositions and catch-all propositions, namely, that we know "what it says" of the former, but not of the latter. This distinction is typically known as the distinction between *transparent* propositions (of which we know the content) and *opaque* propositions (of which we don't). So, whereas *poems* is transparent, a catch-all proposition such as *another genre* is opaque. While this may initially seem plausible, an old insight from Frege (1892) should give us pause. Frege's insight is that any particular state of affairs may be denoted in various ways, and that epistemic contexts are *intensional*, which is to say that the truth-value of epistemic attitude reports are sensitive to the way in which states of affairs are designated.[21] To illustrate how that is relevant to the problem at hand, let us borrow an example from Mahtani (2017). A physician is treating two patients, whom she has nicknamed "Grumpy" and "Whiney" after their personality traits. She does not know which of Grumpy and Whiney has disease X and which has disease Y. Which of the designator pairs, "patient with X" and "patient with Y", or "Grumpy" and "Whiney", is transparent and which is opaque? This is a difficult question. On the one hand, she can visually distinguish between Grumpy and Whiney. So, if she is in their shared room talking to them, this will presumably be the most transparent way she has of referring to them. She will know what the proposition expressing that Grumpy has white hair means, and she will be able to settle it on the basis of observation. However, she will not be able to settle the proposition expressing that the patient with disease X has white hair: after all, there is a sense in which she

[21] For a discussion of Frege's puzzle in Bayesian epistemology, see Chalmers (2011).

does not know who this patient is. But this might change as the context changes. Suppose that the physician can distinguish between the patient with X and the patient with Y on the basis of blood tests. So, if she attends a meeting to discuss the patients' cases, referring to them by the name of their disease will presumably be the most transparent option. Furthermore, she might be able to settle propositions such as the one expressing that the patient with disease X is anaemic, but not be able to settle whether Whiney is anaemic. What this case suggests is that the distinction between transparent and opaque propositions might not be as straightforward as initially thought.

But even supposing that we can make sense of the opaque/transparent distinction, it is unclear that it can do the job of delineating actively considered from catch-all propositions. Suppose that a theoretical physicist tells me that there are two positive candidates for a theory of fundamental physics: super-symmetry theory and quantum gravity theory. In this situation and knowing that these two theories do not exhaust the space of possibilities, I might have credences in *super-symmetry*, *quantum gravity*, and *other*. But, having no training in theoretical physics, I do not know what super-symmetry theory or quantum gravity theory say. Faced with a statement of either of them, I would not be able to recognise it as such. I have no understanding of what it would take for either of them to be true. In other words, they are as opaque to me as *other*. But I take it that, in this example, *super-symmetry* and *quantum gravity* are supposed to be actively considered, and *other* is supposed to be a catch-all proposition. What this shows is that what was supposed to be the distinguishing mark of the catch-all proposition—opacity—may not characteristic of it after all.

But suppose that my objector manages to make precise sense of the considered/catch-all distinction. He still needs to argue that (2) this distinction entails that considered propositions are more easily assigned a justified credence than catch-all propositions. I want to suggest that this is false. We must first note that, as Maher (1995) remarks, "we can have subjective probabilities for [opaque] propositions. For example, I might think it unlikely that the next theory to be formulated will be

simpler than any existing theory" (p. 105). Similarly, our book-lover might have a high credence that the book she will receive is not a novel or a poetry collection. But more importantly, our credences about opaque propositions can seem justified, maybe even more justified than credences about transparent propositions. For instance, even though our book lover may know the content of *poems* and *novels*, she may have absolutely no idea how likely they are relative to one another. By contrast, if she has a high credence that the person offering her the book is an experimental literature enthusiast, she might have good reason to believe that the book is likely not to be as banal as a mere novel or poetry collection. Similarly, I may not know what the next paradigm of fundamental physics will be, but I can still have a high credence that it will be true, higher in fact than my credence that any of the current working theories are true. In light of these examples, we need to reassess the reliability of our initial intuition that the catch-all is supposedly intractable.

So, let me conclude on the widespread argument for the defectiveness view. The reason behind the defectiveness view was meant to be that there is a property that only catch-all propositions have, which is (at least akin to) the property of being opaque, and which makes it uniquely hard to assign it a justified credence. Contrary to this popular claim, I have shown that the transparent/opaque distinction is much less well defined than it is taken to be, that it is at best unclear that only catch-all propositions are opaque, and that whatever it is that makes catch-all propositions supposedly special, it does not make a credence in them any harder to justify. The case for the defectiveness claim is seriously weakened.

Let me now turn to my positive argument against the defectiveness claim. This is a disjunctive argument, composed of a poric sub-argument and a telic sub-argument. Let us begin with the former. The fact that the Bayesian agent has a perfect capacity to reason entails that there is a distinction, from her perspective, between two kinds of propositions: those that she can settle by reason alone, and those that she cannot. There is a distinction in kind between, say, $\textit{novel} \vee \neg \textit{novel}$ on the one hand and *poems* on the other, and this distinction originates in the

epistemic states she has the means to rationally achieve: her means warrant a credence of 1 in the former, but not in the latter. Furthermore, the fact that she has a perfect capacity to observe entails that there is a distinction within those propositions that she cannot settle by reason alone, between another two types of propositions: those that her observations suffice to settle, and those that they do not. So assume that she has observed that she will get a book of fiction, but does not know what kind. There is a distinction for her between *fiction* and *poems*: her evidence allows her to settle the former—that is, to become certain in the truth of the former—but not the latter. Thus the means of inquiry the agent has bring about a tripartite distinction between propositions: the non-empirical propositions, the empirical propositions about the observed, and the empirical propositions about the unobserved.

It is interesting here to note that the agent's means have singled out a particular type of proposition, the empirical propositions about the unobserved, as those which her means do not suffice to settle. In "Bayesianism and the Problem of Induction", I argued that this is the problem of induction, properly understood. How does this relate to the defectiveness claim? Well, those who take the claim to be true believe that it is easier for the agent to have justified credences in, say, *poems*, than it is to have justified credences in the catch-all proposition. But we are now in a position to see why, on the poric approach, this cannot be true. For what means do they have to warrant a particular credence in the former, that does not warrant a particular credence in the latter? Reason alone cannot warrant a particular credence in either, because both propositions have empirical content. Observation cannot warrant a particular credence in them either, because both are about the unobserved. So, we are in a bind when it comes to assigning justified credences to them; but more importantly for our present purposes, we are in the same bind. The question that Bradley put so well holds for both propositions in the very same way: "on what basis are we to determine [...] what probability we should assign to the[m]?". There is no normatively relevant difference between "actively considered" propositions and catch-all propositions in the Bayesian framework.

So far, I have presented a poric argument against the defectiveness claim, by arguing that *qua* propositions about the unobserved, propositions like *poems*, *novels*, and *other* were equally difficult to assign a justified credence to.[22] But one might adopt a means-ends conception of epistemology (as we are assuming), all the while rejecting the poric approach to justification. For instance, one could adopt a telic approach to justification instead. (I have raised doubts in "Why Subjectivism?" as to whether this is a good alternative to the poric approach—but let us assume that it is, for the sake of argument.) The most successful telic strategy for justifying credences about the unobserved is the appeal to the *indifference principle*.[23] According to this principle, any set of propositions about the unobserved that are mutually inconsistent and exhaustive must be assigned an equal credence. But this entails that it is equally easy to determine how confident to be in *poems*, *novel*, and *other*. So, on the prominent telic approach too, we have an argument against the defectiveness claim.

Let me wrap up this section. I began by considering the argument in favour of the defectiveness claim, and arguing that it fails at all stages. Then, I proposed a disjunctive argument against the claim: I showed that it is false on the poric view, and false on the best telic view. I conclude that, despite the vast support in its favour, the defectiveness claim should be abandoned.

4.

In §2, I argued against the contrivance claim—the claim according to which the only reason for including a catch-all in agents' algebras is to accommodate awareness growth. In §3, I argued against the defectiveness claim—the claim according to which catch-all propositions are uniquely difficult to assign justified credences

[22] One could reject this conclusion by rejecting the dualism assumption, and arguing that we have (at least) a means of inquiry, beyond reason and observation, which would warrant particular credences in "considered" propositions but not in catch-all propositions. For reasons I have rehearsed throughout this thesis, I am sceptical that this of this strategy's potential to deliver.

[23] See Williamson (2010), Pettigrew (2016a,b).

to. In the first part of this section, I show that these conclusions can be used to build a strong case for the refinement view, and against the expansion view. In the second part, I examine the consequences of adopting the refinement view for the debate on the norms of rationality that govern awareness growth.

As we saw in §2, a less-than-fully-aware agent must always have a catch-all proposition in her algebra. Furthermore, as we saw in §3, there is nothing more wrong with catch-all propositions than with any other kind of proposition. So we can embrace the fact that our book lover's initial algebra must contain *poems*, *novels*, and *other*, and moreover, that upon becoming aware that the book could be of short stories, her subsequent algebra must then contain *poems*, *novels*, *stories*, and *other*. How is this related to the question of whether lateral awareness growth should be modelled by refinement or by expansion? The reason that those who reject the refinement view do so, is because they hold the conjunction of the contrivance claim and the defectiveness claim. That is, they think that there is a problem with catch-all propositions (defectiveness), refuse to include catch-all propositions in an agent's algebra as a result, and therefore reject the refinement view (contrivance). So, as we reject the defectiveness and the contrivance claims, the reason to uphold the expansion view disappears. But it does not follow that the refinement view is true: the rejection of the contrivance and defectiveness claims does not entail the refutation of the expansion view.

Nonetheless, I think that the most natural interpretation of the Bayesian formalism supports the refinement view over the expansion view, given the inclusion of a catch-all proposition. In §1, I drew the distinction between the two views with a diagram, but a more precise definition is needed. The hallmark of the refinement view is that the later algebra $\mathcal{A}'$ is a fine-graining of the earlier algebra $\mathcal{A}$. What this means is that for any proposition $A_i \in \mathcal{A}$, it must be that $A_i \in \mathcal{A}'$; and there are some propositions $A_j \in \mathcal{A}'$ such that $A_j \notin \mathcal{A}$. So, not only does $\mathcal{A}'$ contain strictly more propositions that $\mathcal{A}$ does (we are, after all, modelling awareness *growth*), but—and this is the hallmark of the refinement view—all the propositions in $\mathcal{A}$ must also be in $\mathcal{A}'$. In other words, $\mathcal{A} \subset \mathcal{A}'$. By contrast, on the expansion

view, $\mathcal{A}'$ must also contain strictly more propositions that $\mathcal{A}$, but there are no further requirements.

At this stage, one might object. Is it not the case that the agent who undergoes awareness growth by expansion also retains a credence in each proposition she considers at the outset? When an agent goes from considering *novels*, *poems*, and *other* to considering *novels*, *poems*, *other*, and *stories*, does she not continue to entertain the initial three propositions? To see that she does not, one must attend to the fact that the propositions entertained by the agent are mandated to form a Boolean algebra. This entails, among other things, that *other* in the agent's initial algebra means *neither poems nor novel*. By contrast, the proposition that one might call "other" in the later algebra does not mean *neither poems nor novel*, but instead means *neither poems nor novels nor stories*. Furthermore, this holds of all propositions. So *poems* in the initial algebra means *neither novel nor (not a novel or a poem)*, whereas "poems" in the later algebra means *neither novel nor stories nor (not novel, poems, or stories)*. As we can see, none of the propositions in the initial algebra are retained in the later algebra, on the expansion view. Thus the debate between proponents of the refinement view and proponents of the expansion view does indeed concern whether an agent ceases to consider any propositions as she undergoes awareness growth.

In what follows, I shall argue (with the help of an analogy) for the refinement view. Consider an agent who initially entertains the propositions that the book of fiction might be a novel or a book of poems, and subsequently realises that, if it is a book of poems, it might be a book of prose poems or a book of poems in verse. (We can falsely assume that this is an exhaustive disjunction for simplicity.) This agent undergoes what I have called internal awareness growth. As I asserted in §1, everyone agrees that this kind of awareness growth should be modelled by refinement. The agent's initial algebra $\mathcal{A}$ has for atomic propositions *poems*, *novel*, and *other*; and her later algebra $\mathcal{A}'$ also has *novel* and *other*, but instead of *poems*, it has *prose* and *verse*. (The agent continues to consider the proposition *poems*, not as an atomic proposition, but as the disjunction of *prose* and *verse*.) The thought

behind this modelling choice is rarely articulated but it is easy enough to discern: what happens as the agent undergoes internal awareness growth is not that she gets an entirely new map—an entirely different representation of the world—rather, it is that she comes to appreciate that her map was insufficiently precise in a place, and modifies this place as a result, retaining the rest of the map.

This line of thought can be transferred directly to lateral awareness growth. Consider the agent who entertains *novel*, *poems*, and *other* at an early stage, and who later comes to entertain *stories*. A natural reading of this situation is that, just like the previous agent realised that poems could be made of verse or prose, this agent realises that fiction that is neither a novel nor a poetry collection could be either a book of short stories or something else still. So, the propositions *novel*, *poems*, and *other* are all still contained in her algebra (she has not ceased entertaining any of them), but the latter is no longer an atomic proposition; instead, it is the disjunction of two newly formulated propositions: *stories* and *other*′. So, on the refinement view, awareness growth consists in coming to appreciate a distinction one was previously unable to make. Whereas at the earlier time, the agent cannot discriminate between literary genres besides the novel and the poem—all non-novel non-poetic genres are lumped together for her—she comes to discern a distinction among these genres, between short stories on the one hand, and yet-other genres on the other.

This suggests two reasons to adopt the refinement view for lateral awareness growth. The first is that the agent who undergoes awareness growth does not cease to entertain any proposition (or, does not find that the truth-conditions for the possibilities that she entertains have changed). Instead, she becomes capable of drawing a new distinction. This is captured by the refinement view, and not by the expansion view. The second is that, if we adopt the refinement view for internal awareness growth (as it is uncontroversial that we should), we should also adopt it for lateral awareness growth. The difference between internal and lateral awareness growth is that the agent comes to discern two types of poetry in the former, and two types of non-novel-non-poetry in the latter. But given my arguments in §3

that there is no normatively relevant difference between "regular" propositions and catch-all propositions, and given that appealing to such a difference would be the only way to create a rift between internal and lateral forms of awareness growth, we must conclude that the two ought to be modelled identically: by refinement. So, to conclude on the first part of this section, given that less-than-fully-aware agents are mandated to include a catch-all proposition in their algebra, the refinement view is the most plausible account of lateral awareness growth.

I want to finish this paper by considering the question, not of awareness growth *per se*, but of belief-change as a result of awareness growth. Take our book loving agent. How should she modify her credences upon becoming aware that the book might be a book a short stories? (For simplicity, let us assume that the possible genres of the book she will receive is the only thing that she considers.) Since, by assumption, she has acquired no evidence, it seems that she ought to retain her previous credences in all the propositions that she entertains across the growth in awareness. She should remain as confident in *novel*, *poems*, and *other* as she was before.[24] But how confident should she be in the relative plausibility of *stories* and *other′*—that is, in the newly considered propositions? Drawing on my arguments in "Bayesianism and the Problem of Induction", I want to suggest that this question is an instance of the problem of induction. Indeed, by assumption, she has no

[24] This thought has been captured in the literature on awareness growth by a norm called *reverse Bayesianism*. It was explicitly proposed by Karni and Vierø (2013, 2015); endorsed by Henderson et al. (2010), Wenmackers and Romeijn (2016), and Bradley (2017a); and rejected by, for instance, Carr (2015), Mahtani (forthcoming), and Steele and Stefánsson (ms). A number of complaints about reverse Bayesianism feature cases where the agent's awareness growth is not a case of pure lateral awareness growth. For instance, Steele and Stefánsson consider a case where an agent comes to consider that a film he thought would be either in French or in German might be either low- or high-brow. But because the agent had antecedent beliefs that French things were more likely to be high-brow than German things, the growth in awareness in their case entails a change in how likely the agent takes the initial propositions to be. There is a lot to say about cases like this—much more than I can hope to say in the conclusion of this paper.

evidence about whether *stories* or *other′*. So, what credences would she be justified in ascribing to these propositions? This question clearly does not have an easy answer—but it is no harder to answer than the question which pervades Bayesian epistemology, of what credence to assign to any proposition about the unobserved. And just like with any other instance of this question, there are several options. The subjectivist will insist that no credal distribution over these two propositions is more justified than any other. The objectivist will disagree, and insist that the agent ought to be equally confident in *stories* than in *other′*. My aim is not to adjudicate between different putative responses to the problem of induction here[25] (in fact, I have already argued for a particular one in "Why Subjectivism?"), but merely to identify the heart of problem of belief-change with awareness growth, and to show that, in Bayesian epistemology, the problem of induction arises again and again.

[25] The objectivist will presumably have instructed the agent, prior to her undergoing awareness growth, to assign credence $1/3$ to *novels*, to *poems*, and to *other*. Then, assuming that she ought not change her credences in propositions upon awareness growth—after all, she gains no evidence—the objectivist will instruct her to assign credence $1/6$ to both *stories* and *other′*. But why should she have a lower credence in say *stories* than in *poems*, given that her evidence is entirely silent between the two? This seems to be a problem by the objectivist's own lights. Of course, this is a famous problem for objectivism. But if—as I think it is—awareness growth is a widespread and frequent phenomenon, the problem might arise more frequently than is typically appreciated.

4. Objectivity and the Method of Arbitrary Functions

Abstract. There is widespread excitement in the literature about the method of arbitrary functions: many take it to show that it is from the dynamics of systems that the objectivity of probabilities emerge. In this paper, I differentiate three ways in which a probability function might be objective, and I argue that the method of arbitrary functions cannot help us show that dynamics objectivise probabilities in any of these senses.

0. Introduction

Probabilities are ubiquitous in a number of special sciences, including statistical mechanics, population genetics, and many social sciences. Within these sciences, they seem to play a crucial role in the explanation and the prediction of empirical phenomena, and to provide constraints on what we ought to believe about these phenomena. This has motivated many to think that they must be *objective*, in some sense to be made precise. Are they, and if so, in what sense? There is a vast literature on these topics, a subset of which concerns what is known as the *method of arbitrary functions*.[1] This method will be my central concern in this paper.

The method of arbitrary functions has attracted a lot of attention, because it is widely believed to provide a reason for thinking that (at least some of) the

[1] The method originates with von Kries (1886) and Poincaré (1896). For a history of the method, see von Plato (1983).

probabilities in the special sciences are indeed objective, in virtue of the dynamics of the systems under study. More precisely, many take it to show that the dynamics of a system can serve to *objectivise* an underlying probability function; or in other words, that the objectivity of the probabilities can emerge from the dynamics of the system. Strevens (2011) for instance claims that the method allows us to

> ...find a basis for an outcome's [objective] probability in the properties of the physical dynamics that produce it. (p. 340)

According to Myrvold (2016), the key idea of the method is that

> ...a probability distribution over initial conditions is transformed, via the dynamical evolution of the system, into a probability distribution [which is] an epistemic chance. (p. 592)

For Gallow (forthcoming),

> ...the chance distribution over outcomes is the result of time-evolving a suitable probability distribution along a suitable dynamical equation. [...] In a slogan: chance is a suitable probability distribution filtered through suitable dynamics.

All these authors agree: the method of arbitrary functions shows us that the objectivity (or chanciness) of a probability function emerges out of the system's dynamics. My aim in this paper is to argue that the method of arbitrary functions can do no such thing.

My argument begins by differentiating three ways in which a probability function might be thought to be objective. A probability function might be ontically (as opposed to epistemically) interpreted, if it represents a mind-independent phenomenon. If it is interpreted as a credence function, it might be objectively (as opposed to subjectively) evaluable, if disagreement about its values entail fault. Finally, it might be high-level robust (as opposed to chaotic) if the values it assigns to high-level phenomena do not depend much on the values it assigns to low-level

phenomena.[2] In this paper, I show that authors in the literature on the method of arbitrary functions have sought to use the method to show that the dynamics of a system can establish the objectivity of probability functions in each of these senses. I then show that this is a mistake: the method does not show that the objectivity of probability functions stems from the system's dynamics, in any of the senses of objectivity that can be found in the relevant literature.

Let me make my thesis more precise. What I will argue is that, to the extent that the aforementioned probabilities are objective, they are not so in virtue of dynamics *qua* dynamics. Thus my argument will be consistent with the claim that the probabilities in question are objective; and it will even be consistent with the claim that the method of arbitrary functions can shed some light as to why they are objective (though for this latter claim, not in all senses of *objective*). But, I will show, this has nothing in particular to do with dynamics, rather at best, dynamics are an instance of a much more general way in which the probabilities in question could be objectivised. Thus my claim will be that the authors are wrong to think that the basis for the objectivity of probabilities lies in the dynamics (or mechanics) of the systems of study. It follows that, if authors are to argue that they are indeed objective, in any of the senses I have identified, they cannot appeal to the dynamics of the system in question.

My strategy is as follows. I begin by outlining the method in some detail (§1). Then, I take each of the senses of objectivity in turn. My strategy is the same for each of them. I begin by giving a fuller account of the sense of objectivity at hand. Then, I show that at least some authors in the literature believe the method of arbitrary functions to show that it is the system's dynamics that establish the

[2] There is also another sense of objectivity, according to which something is objective if it is perspective-invariant; if it is "the same from all angles". This sense of objectivity has been discussed in the literature on objectivity in general, as well as in the literature on the objectivity of probabilities in the form of the reference class problem. However, as far as I'm aware, it has not been discussed in the literature on arbitrary functions. Nonetheless I discuss it briefly in §4, where I show that the method does not establish the objectivity of probabilities in the sense of addressing the reference class problem.

objectivity of probabilities in that sense. Finally, I argue that they are wrong. To the extent that probabilities in the sciences are ontically interpreted (§2), objectively evaluable (§3), and high-level robust (§4), they are not so in virtue of the dynamics of the system. I conclude the paper by noting that, even though the method is inert when it comes to the objectivisation of probabilities through dynamics, it nonetheless has philosophical interest: I explain how (§5).

1.

In this section, I outline the method of arbitrary functions. Suppose that you are faced with a wheel, painted with alternating black and red wedges of equal size, equipped with a stationary pointer. The wheel may be spun, and allowed to come at a stop with the pointer indicating either red or black. Suppose further that the mechanics of the wheel are deterministic, and that the outcome of a given trial is fixed by a single parameter, the initial speed of the wheel.

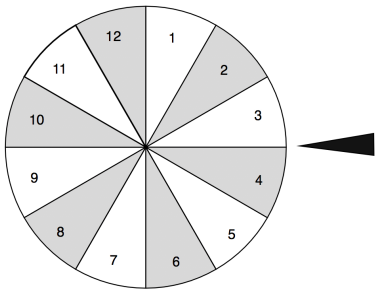

1. A roulette wheel

Let us represent this setup more precisely. Let the set X represent the set of all initial speeds, the set Y be the set of all the wedges on the roulette wheel,[3] and the set Z be the set of all outcomes of interest—here, black and red. Now, let us define a surjective function $f : X \to Y$, which maps every initial speed to its associated outcome on the roulette wheel.[4] Next, we define a surjective function

[3] Alternatively, we could represent the setup with a set representing the wheel's possible angles of rotation.

[4] This map is a function because the dynamics of the wheel are deterministic: given that

$g : Y \to Z$, which maps every wedge of the wheel to its colour. Finally, we define a surjective function $h = g \circ f : X \to Z$, which maps initial speeds to the colour of the wedge that it yields. Note that these functions allow each subset of Y and Z to be associated with a subset of X, namely its preimage. For instance, f associates $\{y_1\}$, the first wedge, with the set of all initial speeds that yield a landing on the first wedge, which we can call X_1.

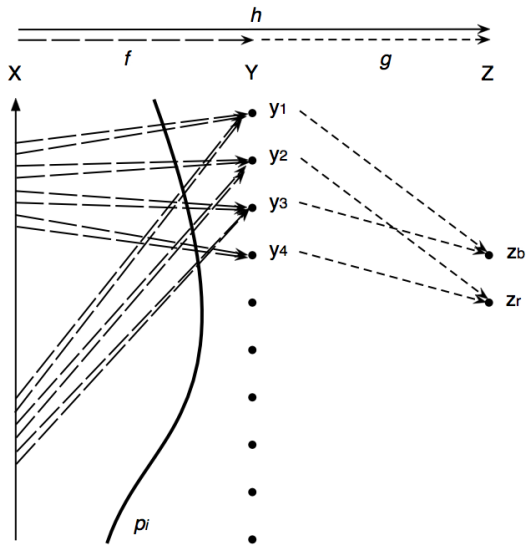

2. The setup

Now, we can introduce probabilities. Let us define a probability function p_i on (a σ-algebra of) X, to give us the probabilities of (sets of) initial speeds. (I leave questions of interpretation to §2.) Note that this probability function will give us the probability of any measurable subset of X. So for instance, it will give us the probability of X_1, the set of all initial speeds that yield a landing on the first wedge;

---

the wheel is spun at a particular initial speed, there is only one wedge that it could land on. This suggests that, in the event that someone wanted to extend the method of arbitrary functions to non-deterministic sciences, such as quantum mechanics, some amendments would have to be made. I leave this aside, though: along the other authors writing about the method of arbitrary functions, I concentrate on the probabilities that feature in the deterministic sciences.

and the probability of X_B, the set of all initial speeds that yield a landing on a black wedge.[5] Note that, strictly speaking, this does not give us the probability of landing on a specific wedge, or the probability of landing on a wedge of a specific colour: it gives us the probability of the wheel being spun with an initial speed such that it will land on a specific wedge or colour. Finally, note that strictly speaking, the probability function p_i does not ascribe values to individual initial speeds: it is its associated probability density function that does.[6] However, for economy of expression, I will obscure this distinction in the rest of this paper and sometimes talk of p_i as taking initial speeds as its argument—talking for instance of "the probability of an initial speed" or of "the probability function over initial speeds".[7]

We arrive at the method. The basic idea is that, from two relatively weak assumptions, we can derive that the probability of X_B is roughly equal to that of X_R. The first assumption **(a)** concerns the mechanics of the system; that is, the shape of the functions f, g, and h. It can be divided into two parts. **($\mathbf{a}_1$)** We assume that f^{-1}, the inverse image of f, partitions X into intervals such that velocities in adjacent intervals are mapped to adjacent wedges: the resulting partition is what Butterfield (2011) calls "filamentous" (p. 1083). Roughly, if you slightly increase the initial speed by which the wheel is spun, the wheel will land on the following wedge. **($\mathbf{a}_2$)** Furthermore, in sufficiently small (but not too small) regions of initial speeds, the relative size of the set of initial speeds mapped by h onto the red outcome with respect to that of the set mapped by h onto the black outcome is roughly constant: the system has a property which Strevens (2011) calls "microconstancy" (p. 346).

[5] We are assuming that f and h are random variables; that is, that the preimages of elements of the algebras over Y and Z are measurable in the algebra over X.

[6] Indeed the probability density function is defined on elements of X, that is, over X; and its associated probability function p_i is defined on measurable subsets of X, that is, over a σ-algebra of X. So p_i takes *sets* of initial speeds as its argument, not initial speeds themselves.

[7] In addition, I have labelled the probability density function in figures **2**–**4**, "p_i". Again this is not strictly correct because elsewhere I take "p_i" to refer to its associated probability function. Nonetheless I remain loose for economy of expression.

In our roulette wheel, the ratio is $1/2$. The second assumption **(b)** concerns the probability function p_i over sets of initial speeds. We assume that it is such that its associated density function does not vary too quickly, such that two very similar initial speeds get roughly the same probability density.

If these two assumptions are satisfied, it can be shown that the probability of X_B is roughly equal to that of X_R.[8] This is illustrated by figure **3** below. Given a setup like that of figure **2**, which respects assumption **(a)**, take any function (that is, take an ***arbitrary function***) p_i that respects assumption **(b)**, and it follows that $p_i(X_B) \approx p_i(X_R)$. This result is what is usually referred to as the "method of arbitrary functions". (I will sometimes speak as though $p_i(X_B)$ and $p_i(X_R)$ were equal, as opposed to roughly equal. This is merely for economy of expression—they are not.)

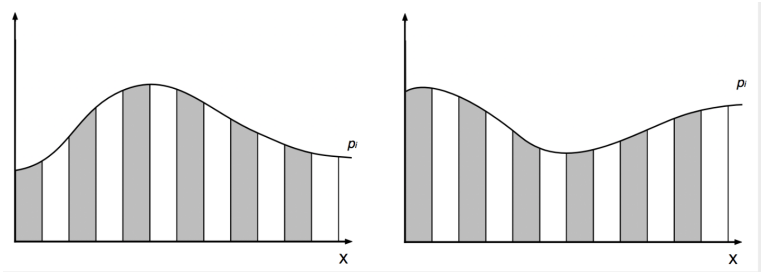

3. Two different probability functions over $\mathcal{A}$.

You can see that the two assumptions **(a)** and **(b)** pull in opposite directions. The more "filamentous" the partition of X is, the more "wiggly" p_i can be (see figure **4.1**). Conversely, the less "wiggly" p_i is, the less "filamentous" the partition needs to be (see figure **4.2**).

[8] The initial proof appears in Poincaré (1896). See Butterfield (2011) for an overview of the technical results, and Engel (1992) for a more in-depth (and difficult!) exposition.

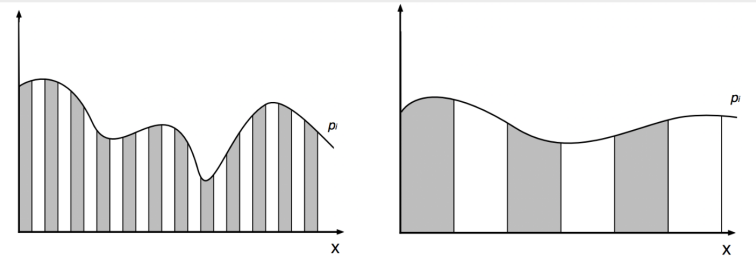

4. Assumptions **(a)** and **(b)** pull in different directions.

With a grasp of the method of arbitrary functions, we are in a position to ask: what is the philosophical significance of this result? In particular, can this result help in establishing the objectivity of probabilities in the special sciences?

2.

In this section, I consider whether the method of arbitrary functions can help establish that some probabilities in the sciences are objective, in the sense that they are *ontically interpreted*.

What is objectivity in this sense? I begin by spelling out what an interpretation is. Imagine that there are two people in a room. We might represent these two people using a set, $A = \{a_1, a_2\}$, each element standing for a person. Now, suppose that we are interested in the height of these people in centimetres. We might represent that using a function $h : A \to \mathbb{R}$, which assigns a real number to each element of A, giving the height in centimetres of each person represented as an element of A. In that case, we say that the function h *represents* the height of the two persons. Conversely, imagine that one is presented with a function $h : A \to \mathbb{R}$. On its own, h is a purely mathematical object, and does not describe anything in the world. However, we might *interpret* that function. What that means, is that we specify what the elements of the set A on which h is defined represent, and what h says of them. Interpreting h as a height function thus requires identifying each element of the set A with a person in the room, and understanding h as giving the height in centimetres of each of those people. Understood in this way, representation and interpretation are *duals*. Give me a fact and I can represent it

using a function, give me a function and I can interpret it as expressing a fact.[9]

We can divide interpretations of functions in two groups: the *epistemic* ones, and the *ontic* ones. We say that an interpretation is epistemic just in case the function under that interpretation describes an agent's epistemic, intentional attitude.[10] We say by contrast that it is ontic if a function under it describes something non-epistemic; that is, mind-independent. A snappy way to think about this distinction is that ontically interpreted functions are about the external world, whereas epistemically interpreted functions are (at least partly) about the mind. To illustrate, take the function $h : A \to \mathbb{R}$. The height interpretation of this function described above is ontic: the height of the people in the room is unconnected to anyone's epistemic attitudes. By contrast, an interpretation of the function h as the height that some agent Alice believes the people in the room to have is epistemic.

Like other functions, probability functions can be used to describe the world, and as such, can be interpreted. The question of interpretation is that of what these functions represent in the world. So, take a probability function $p : \mathcal{A} \to \mathbb{R}$, where $\mathcal{A} = \{\Omega, \{H\}, \{T\}, \emptyset\}$ and $p(\{H\}) = p(\{T\}) = 1/2$. What it means to interpret this function, is to specify what the elements of $\mathcal{A}$ represent, and what p says of them. One way to interpret this function is as follows. A coin has been flipped ten times. The elements of the algebra correspond to the different possible outcomes of the coin flip: coin lands heads, and coin lands tails (and the trivial events). The probability function p describes the proportion (that is, the finite frequency) of these outcomes having obtained. Namely, it says that the coin has landed half the time on heads and half the time on tails in this sequence of flips.

[9] This formulation is awkward, but the hope however is that this snappy phrase, problematic as it may be, gives the reader some initial insight into what I take interpretation to be.

[10] Technically, there may be interpretations of probability that represent intentional attitudes that are not epistemic. For instance, we can imagine a probability function interpreted as giving an agent's degree of desire that a proposition be true. Thus not all possible mind-dependent interpretations are subjective. But I don't think that these are relevant to the matter at hand; so I leave them aside.

Here, p has been ontically interpreted. Another way to interpret p is as follows. A coin is about to be flipped, and Alice is wondering how it will land. The elements of the algebra correspond to the propositions that the coin will land heads, and that the coin will land tails (and the trivial propositions). The probability function p describes the confidence Alice has in the truth of each such proposition: it says that she is equally confident that the coin will land heads and tails. We say that p represents her *degrees of belief*, or *credences*, in these propositions; so here, it has been epistemically interpreted.

It is customary in the literature to call the ontic interpretations of probability "objective",[11] and the epistemic ones "subjective".[12] Many think that probabilities in the sciences should be objective in this sense: the (physical) sciences describe the world and not our minds. Moreover, these probabilities play an explanatory role: they are used for instance in our best statistical mechanical explanation of why milk dissolves in coffee. But they could not play this role, the thought goes, if they described the mind and not the (mind-independent) system at hand.[13] Can the method of arbitrary functions help establish that probabilities are objective in this sense? Many in the literature think that it can. Some claim this directly. Abrams (2012) purports to use the method of arbitrary functions to "outline a new interpretation of probability, one which is objective in the sense that the resulting probabilities are constituted only by facts about states of the world, without regard to epistemic factors such as belief or justification" (p. 344). Rosenthal (2012) argues that the method can "yield an objective interpretation of probability in the sense of providing truth conditions for probability statements that are independent of our state of mind and our state of information" (p. 225). Werndl (2015) argues that

[11] A number of different objective interpretations of probability are discussed in the literature, including finite frequency (as above), hypothetical frequency, propensity, and indeterministic chance. For an overview, see Hájek (2019). Since my point applies to all interpretations indiscriminately, I will not discuss their particular characteristics in any detail.

[12] The main subjective interpretation is the Bayesian one, described above.

[13] For a forceful defence of this point, see Albert (2000).

the method "shows how stable ontic probabilities (i.e., probabilities that are real features of the world) can arise out of deterministic equations" (p. 223).

But it is interesting to note that many other authors are implicitly committed to this view. Central to the debate about the compatibility of determinism and objective probability[14] is a sceptical argument; roughly: what can it mean to say that the objective probability of a coin landing heads is $1/2$, if it is already determined by the laws of nature that it will? The sense of objectivity that is relevant here is that of ontic interpretation. Indeed, the question is that of what these probabilities could represent that is not epistemic.[15] So, to think that progress can been made in the compatibilist camp amounts to thinking that progress can be made in devising an ontic interpretation of probability. And indeed, many mention the method of arbitrary functions in connection with the compatibilism debate, including Butterfield (2011), Frigg (2016), Bradley (2017b), and Hájek (2019). But if I am right that the method can provide no help when it comes to ontically interpreting probabilities, these authors are mistaken in thinking that there is a connection between the method and the compatibilism debate.

Am I right, then, in thinking that the method can provide no such help? Or could it in fact at least make progress on establishing an ontic interpretation of the claim that red and black outcomes of a wheel spin are equally likely? To start, note that the method is really a mathematical theorem. So, in order for it to say anything about the world (and thus have philosophical significance), the mathematical objects that comprise it must be interpreted. The functions f, g, and h are interpreted straightforwardly, as representing the dynamics and structure of the roulette wheel. But what about p_i? I will not answer this question here. Rather, I will simply note that whichever interpretation is supplied to p_i is exogenous to the method of arbitrary functions; and furthermore that p_i is the only probability function mentioned throughout. It follows that the method is completely inert, as far as interpretation is concerned: it cannot help establish the objectivity of the

[14] For an overview, see Frigg (2016).

[15] Note for instance that there is no tension between determinism and objective Bayesianism (see §3 for a discussion of this sense of objectivity).

equiprobability of red and black outcomes. This is all the more the case given that the method does not even make claims about the probabilities of red and black outcomes themselves: instead, it makes claims about the probabilities of sets of initial speeds that will yield these outcomes. To illustrate, let us assume that p_i is interpreted as giving frequencies. The method then says that, given **(a)** the mechanics and structure of the wheel, and **(b)** that the frequencies of initial speeds do not vary too quickly, the sets of initial speeds that correspond to red and black outcomes occur with roughly equal frequency. It should be clear that the method has played no interpretive role.

The three authors cited in §0 believe that, contrary to what I have said so far, the method of arbitrary functions does make claims about the probability of landing on a red wedge; as opposed to merely making claims about the probability of the initial speed being such that the wheel will land on a red wedge. In order for this be true, the method must be supplemented. A new probability function, call it P, must be defined on (a σ-algebra over) Z, ascribing probabilities to the two possible colour outcomes of wheel spins, red and black. The method of arbitrary functions must then be connected to this new function by fixing its values. What we want is for the probability of landing on a black wedge to be equal to the probability of the wheel having been spun with a black-yielding initial speed; that is $P(z_B) = p_i(X_B)$, and similarly for the red outcome. We now have a second probability function, whose interpretation is not exogenously fixed. Could the method so-supplemented then help with interpretive projects?

There are two options. The first is that P is given the same interpretation as p_i; more precisely, P is stipulated to say of z_B what p_i says of X_B. So, on this option and assuming that p_i was interpreted as giving frequencies, the (supplemented) method of arbitrary functions tells us that, if the frequencies of initial speeds do not vary too quickly, then the two colour outcomes occur with equal frequency. But the method has done no interpretative work here: the interpretation of p_i was supplied exogenously, and the interpretation of P was simply stipulated to be identical. This leaves us with the second option, according to which P is not

stipulated to have the same interpretation as p_i. But what interpretation does it then have? The method has fixed its value but has said nothing of its interpretation; which would again have to be supplied exogenously. So, either way, the method of arbitrary function has been of no interpretative help.

In sum, the method of arbitrary functions cannot help establish the objectivity in this first sense of probability distributions over outcomes; that is, it cannot help establish that they can be ontically interpreted. It follows, as I explained earlier in this section, that it is irrelevant to the compatibilism debate.

3.

In this section, I consider whether the method of arbitrary functions can help establish that some probabilities in the sciences are objective, in the sense that they are ***objectively evaluable***, on the basis of the system's dynamics.

What is objectivity in this sense? Consider two agents, Beth and Carl. They are both considering the proposition that a particular spin of our roulette wheel will yield a red outcome, and they possess the same evidence. Suppose that they disagree, which is to say that they have incompatible credences in this proposition: Beth's is $1/3$, and Carl's is $1/2$.[16] Can they simply agree to disagree, or must it instead be that at least one of them is wrong? In particular, is it the case that Carl is right, that $1/2$ is in some sense the objectively correct credence to have in this proposition?

This question has been thoroughly debated. Some think that there is no single objectively correct credence to have in this proposition; among them are the ***Radical Subjective Bayesians***, according to whom every credence function over an al-

[16] I am assuming throughout that credences are probabilistic: I have defined them as epistemically interpreted probability functions. I am also assuming of course that the agent has no direct observational evidence about the outcome of the spin; so these credences may be thought of as priors. Finally, I am assuming that the agent has precise (as opposed to imprecise) credences: we are dealing with a single probability function, and not a set thereof.

gebra of propositions is as good as any other.[17] Others disagree. This includes the *Objective Bayesians*, who think that, given a fixed body of evidence, there is a unique objectively correct credence function—which one is a matter of dispute, though a Principle of Indifference, recommending a uniform distribution over elementary propositions, is often invoked.[18] It also includes the proponents of bridge principles between objectively interpreted probabilities and credences, according to whom credences about objectively probabilistic propositions should correspond to the objective probabilities of the events, though credences in other propositions are unconstrained. The classic chance-credence bridge principle is Lewis' Principal Principle (1980); and the classic frequency-credence bridge principle is proposed by van Fraassen (1983). So according to the proponent of the Principal Principle, all credences in chancy propositions are objectively evaluable, and all credences in non-chancy propositions are only subjectively evaluable.

This suggests a second sense in which a probability function can be objective, based on how much disagreement is allowed between agents. Note firstly that this sense of objectivity applies to epistemic attitudes, and as such, only applies to probabilities interpreted as credences (that is, probabilities that are subjective in the sense discussed in §2). Note secondly that this sense of objectivity is not an all-or-nothing affair. The subjectivist thinks that credences are less objective than what the proponent of the Principal Principle thinks. In turn, the proponent of the Principal Principle thinks that credences are less objective than what the objectivist thinks. Thus we arrive at a more precise definition of this sense of objectivity. Let us say that credences are increasingly *objectively evaluable* as they are subject to increasingly stringent norms of correctness, or in other words, as they allow for a decreasing amount of disagreement. Let us say, derivatively, that norms prescribe certain credences as *objectively correct*, when they dictate that disagreement in these credences entail that at least one agent is wrong. Thus the more objectively correct credences a (true) norm dictates there are, the more objectively evaluable will

[17] This tradition begins with de Finetti (1937).

[18] For a prominent defence of objective Bayesianism, see Williamson (2010).

credences be. Let us illustrate this. According to the objectivist, credences are maximally objectively evaluable: any disagreement between agents entails that at least one has an objectively incorrect credence. According to the proponent of the Principal Principle, credences are somewhat objectively evaluable: disagreement about chancy propositions is disallowed, but other disagreement is allowed. And according to the radical subjectivist, credences are not objectively evaluable at all: all disagreement is allowed.

A clarification must be made about what is meant by a credence being "correct" or a disagreement being "allowed". Let us take the Principal Principle as an example. I have said above that it establishes (if true) that credences towards chance propositions are objectively evaluable, or that disagreement about these propositions is not allowed. This does not entail that agents are irrational if they have the incorrect credence, or if they disagree. For suppose that the chance of some proposition A is $1/2$, but that neither of two agents know this. Suppose furthermore that both of them have perfectly rational credences, in the sense that they have perfectly conditionalised their priors on their available evidence. In this case, the two agents are (internally) rational. But this does not entail that they are correct, given the norm established by the Principal Principle. For according to that norm, if one of them has a credence in A that is not $1/2$, that agent has an objectively incorrect credence; that is, they have a credence that does not correspond to the chance. So, agents cannot simply "agree to disagree": if they have different credences about A, at least one of them has an incorrect credence.[19] Note that we can know that credences are objectively evaluable, without knowing which is objectively correct. If we know that the Principal Principle is true, and that A is a chancy proposition, we know that there is a correct credence to have about A, though we might not know which one, if we do not know what the chance of A is. To conclude on this point: objective evaluability is not a matter of (internal)

[19] Suppose that Jean's credence in A is $1/2$ and that Joan's is $26/50$. It may be that, for all practical purposes, their disagreement does not matter, in the way that it would if their credences were very far apart. But this is not the relevant sense for our purposes: what matters is that they cannot both be strictly correct.

rationality;[20] rather, it is a matter of there being (externally) correct credences.

What would it take for the method of arbitrary functions to show that dynamics objectivise probabilities, in this sense of objective? It would require for the method to show that the system's dynamics provide a reason for thinking that credences are objectively evaluable; that is, that some credal values in some propositions are objectively correct. In the case we have been considering, it would require providing a reason for thinking that Carl has the objectively correct credence. Can the method show that the dynamics of a system provide a reason for $1/2$ being the correct credence to have in the wheel stopping on heads? There are two ways they could do this. The first is that they could help establish that the (mind-independent) chance or frequency of heads is $1/2$. Together with a bridge principle, it would follow that an agent's credence in heads ought to be $1/2$. But there is a sense in which the dynamics themselves would not directly be constraining credences: they would be operating on frequencies/chances, and the bridge principle would do the work of credence constraint. They would not establish a new norm of objectivity, but would instead feed into an already existing norm. I will discuss the role of the method in constraining the values of objectively interpreted probability functions in §4, and the consequences this has for the constraining of credences in §5; but in this section I concentrate on the direct objectivisation of credences.

This brings us to the second option, namely that the method could somehow show that the dynamics of the wheel constrain credences directly. Some have thought that it could do this; that is, that the dynamics of the wheel could generate a reason to think that a credence of $1/2$ in red is objectively correct. This would be attractive, because it would supply a validation, rooted in empirical features of the

[20] Nor is it a matter of practical concerns. There may be some very good reasons for scientists to adopt knowingly incorrect credences, for instance if they are close-enough for all practical purposes and more easily tractable. So it may be that the agent should adopt these credences, but that does not make them correct in the sense I have been describing: they should adopt them, *despite* the fact that they are incorrect, because they have other virtues (such as tractability) which are practically more important in that context.

world, of the intuition shared by many that agents should have this credence in red. The main proponent of this view is Myrvold (2012). He claims that "the dynamics of the system lead all reasonable credences" to a credence of $1/2$ in red (p. 80). Furthermore, he claims that it makes sense to talk, "in cases of disagreement about [the credence to have in red], about one value being more correct than another" (p. 80). So, according to him, the dynamics of the system give us a reason for thinking that there is an objectively correct credence to have in red. In other words, the dynamics of the system objectivise credences, in the sense that they make them (more) objectively evaluable.

How does Myrvold purport to establish this claim? I won't rehearse the precise details of his view here, but simply outline its general structure. This will turn out to be enough to show that he is mistaken in his claims about the power of dynamics to render probabilities objectively evaluable. What then is his view? Firstly, he suggests that we should interpret p_i as an agent's credences in initial speeds, following Savage (1973). Then, he claims that "applying the dynamics of the system to this credence function yields" a probability function P over the set of colour outcomes (p. 79), which gives the "epistemic chance" of these outcomes. He borrows the term "epistemic chance" from Schaffer (2007), who claims that they are subjectively interpreted probabilities which are "objectively informed, and may wear scientific credentials" (p. 137). Thus for Myrvold, the function P represents the credences that an agent ought to have, given the dynamics of the wheel. Finally, he claims that P, together with a principle bridging epistemic chances and credences, can constrain the agent's credences P_i in colour outcomes.[21] Now, there are several question marks surrounding aspects of this view, such as: What does it mean to "apply" dynamics to a credence function? Where does the interpretation of P come from? In particular, where does the normative force of P over an agent's credences come from? How could P have normative force over p_i when P

[21] Myrvold uses different nomenclature to refer to the specific functions; and talks about the epistemic chances, not of colour outcomes, but of the possible angles of orientation of the wheel after it has been spun. I have adapted Myrvold's claims for consistency within this paper, and nothing of philosophical significance hangs on this.

was defined partly in terms of p_i?

We don't need to answer these questions to see that Myrvold's proposal will not succeed in achieving his aim: on his proposal, it is not the dynamics that provide the constraint on credences. For where resides the normativity on his view? It resides in a newly defined function P and on a bridge principle according to which the agent's credences must conform to P. But, by the time this can even be stated, the method of arbitrary functions has already done all its work: indeed as we have seen, what the method shows is that $p_i(X_B)$ and $p_i(X_R)$ are roughly equal given particular assumptions. This must be assumed in order to define the function P, and in turn to outline the bridge principle. So, whatever new reason Myrvold may provide us for having objectively evaluable credences, this reason cannot be the system's dynamics: it must instead be something additional.

This criticism of Myrvold's view suggests that the dynamics of the wheel could only constrain credences if the normativity came from, not an additionally defined function P together with a bridge principle, but from the properties of p_i itself. In what follows, I present such a view, which I show is illuminating in various ways, but does not succeed either in achieving Myrvold's aim. Let me present my view with the help of an example, the *urn case*. Consider a set of six balls, numbered 1 through 6, placed in an urn. Let $F = \{1, 2, 3, 4, 5, 6\}$ be the sample space of possible outcomes of a ball pick, and $\mathcal{F}$ be the power set of F. Let us define a probability function p over $\mathcal{F}$. Suppose that the even-numbered balls are coloured in blue, and the odd-numbered balls are coloured in white. Let us define a random variable $f : F \to F'$, where $F' = \{W, B\}$ is a two-element set that corresponds to white and blue outcomes. Thus f is a coarse-graining map. We have the mathematical resources now to ask, what is the probability of pulling a white ball: what is the value of $p(f^{-1}(\{W\}))$?

Suppose that we interpret p as an agent's credences, and suppose that the agent does not know which unobserved ball is of which colour. More precisely, she doesn't know the shape of f, so that she does not know that $\{1, 3, 5\} = f^{-1}(\{W\})$. In other words, she does not know that "a ball numbered either 1, 3, or 5" and

"a white ball" refer to the same outcome. Because she does not know that these two expressions corefer, she may have different beliefs about them. For instance, if her credences are such that $p(1) = p(3) = p(5) = .1$, then her credence in their disjunction will be $p(\{1, 3, 5\}) = .3$. But supposing that she has also heard from a source that she considers reliable that the white and blue outcomes are equiprobable, her credence in the white outcome might be $p(\{W\}) = .5$. This is no other than a *Frege case*: an agent has two incompatible beliefs about a single state of the world, because she considers this state of the world under two distinct descriptions that she does not realise corefer (Frege, 1892).

The problem raised by Frege cases in general is this: even though it is hard to say what an agent has done wrong in these cases (after all, it doesn't seem to be a failure of rationality that an agent not know which unobserved ball is of which colour), there is a sense in which an agent should not have two distinct credences in one and the same state of affairs. It is notoriously difficult to accommodate for Frege cases in a Bayesian framework:[22] if "a ball numbered 1, 3, or 5" and "a white ball" are represented by a single element of the algebra, it is impossible to represent the agent as having different credences in these possibilities; if they are represented by different elements, then the algebra misrepresents the world. How to accommodate these cases in a Bayesian framework is a problem for another paper, I want only here to remark the following. If an agent has different credences in $\{1, 3, 5\}$ and $\{W\}$, there is some sense in which her credences are objectively incorrect. If two agents had the same credences in $\{1, 3, 5\}$, but had different credences in $\{W\}$, at least one of them would be wrong in the sense described earlier in the section: they could not permissibly "agree to disagree" about the credence to have in $\{W\}$.[23]

[22] See for instance Chalmers (2011).

[23] It follows from this that all cases of agents having non-extremal credences in empirical propositions are involved in a Frege case. Suppose that Julian's credence that it will rain tomorrow is $3/4$. Suppose furthermore that his credence that it will rain on the next rainy day is 1. Assuming that it is indeed true that it will rain tomorrow, the expressions "tomorrow" and "the next rainy day" corefer. So, Julian is making a Frege mistake here:

Let us transpose what I have said about the urn case to the method of arbitrary functions, as presented in §1. Firstly, let us interpret p_i as an agent's credence function over initial speeds. Now, suppose that h and p_i satisfy constraints **(a)** and **(b)**, such that the agent's credence in X_B, namely the set of initial speeds such that the wheel spun at one of those speeds will land on a black wedge, is roughly $1/2$. Suppose furthermore that the agent's credence that the wheel will land on a black wedge is $1/3$, maybe because she has heard from a source that she considers reliable that there are brakes under the red wedges which makes landing on them more likely. The agent is involved in a Frege case: she has different credences in two propositions that are extensionally equivalent, namely the proposition that the wheel is spun such that it will land on a black wedge, and that the wheel will land on a black wedge.[24] But those credences cannot both be correct. If two agents have the same credences about initial speeds, they also ought to have the same credences about colour outcomes. Or in other words, an agent's credences in initial speeds determines the credences to have in colour outcomes, that is, the objectively correct credences. So have we not, through the method of arbitrary functions, found a reason *in the world* for constraining an agent's credences? Has

---

he is assigning two different credences to the same proposition. It is unclear what to do about this fact, especially as it pertains to the specific issue of the method of arbitrary functions. It may however have interesting implications for the more general question of credal disagreement and the aims of credences, but that is another paper.

[24] As discussed above briefly, there are questions about how to represent this agent's epistemic state. If her credences are represented by the function p_i on X, then she must be represented as assigning two different credences to the same proposition X_B. But of course this is impossible, because p_i is a function: by definition, it takes an argument to a single value. There are two options to avoid this. The first is to use two different credence functions to represent the "guises" under which she considers the single proposition. The second is to enlarge her algebra, such that the two extensionally equivalent propositions are represented by two different elements. Which option to pursue is far beyond the scope of this project however. For my current purposes, it suffices to note that agents are susceptible to these Frege cases. How to represent that adequately is a further matter.

the method of arbitrary functions not shown us that the dynamics of a system can serve to objectivise the credences an agent has about that system, in providing a new way to make them objectively evaluable?

I think it has not. In the discussion above, the norm to which we have appealed to establish the objective evaluability of credences makes no reference to the dynamics of the wheel. Instead, the norm is much more general: it prescribes not having different credences in a unique state of affairs. What we did above, was take as a starting point the fact that agents have particular credences in sets of initial speeds (so say, in X_B), and then use this norm to insist that agents should have particular credences in colour outcomes (in the wheel landing on a black wedge). But now we can see that the dynamics of the wheel *qua* dynamics have played no role in this. This is because the dynamics have already played their role by the time the norm can even be defined: the dynamics fix the value of X_B, which is assumed when the norm is stated.

One might object by claiming that fixing the value of X_B is an essential part of establishing that credences towards the black outcome are objectively evaluable. And furthermore the method plays a crucial role in this: it shows that, if agents have credences in individual initial speeds that respect particular constraints, they ought to have credences in sets of initial speeds such that $p_i(X_B) \approx p_i(X_R) \approx 1/2$. Does that not show after all, that the dynamics of the wheel can serve to objectivise credences? It does not. To see this, it suffices to look at the analogue claim within the urn case. The structure of the balls in the urn are such that, if an agent has credences $p(1) = p(3) = p(5) = .1$ in each of these individual balls being drawn, she ought to have credence $p(\{1, 3, 5\}) = .3$ in their disjunction. But the urn case is not dynamical. So, dynamics are not involved in establishing this conclusion. Instead, it is another feature that the urn case and the wheel case have in common, which is involved: the fact that we are interested in a system at various levels of grain. I explore granularity in more depth in the next section.

To conclude on objective evaluability: I have shown in this section that the method of arbitrary functions cannot help establish that the dynamics of a system

can objectivise credences, in the sense of making them more objectively evaluable. To the extent that it does establish the objective evaluability of particular credences, it does so in virtue of a very general norm, according to which it is incorrect to have two conflicting beliefs about a single state of affairs. This suggests that Frege cases in general, and not the method of arbitrary functions in particular, should be studied if one wanted to establish the objective evaluability of probabilities in the sciences. And to the extent that it fixes which credences to have in X_B and X_R, it does so in virtue of comprising a coarse-graining map, not in virtue of representing a dynamical system.

4.

In this section, I consider whether the method of arbitrary functions can help establish that some probabilities in the sciences are objective, in the sense that they are high-level *robust*. What does this mean?

Authors who take this line on the method of arbitrary functions begin with an observation: roulette wheels of the kind we have been considering systematically land on a red wedge and on a black wedge with roughly equal probability. (Let us leave the implied probability function uninterpreted for now.) This raises the question: how can that be? What is it, in the world, that makes this true? The reply given by the authors who take this line is that these probability values are high-level *robust*. What this means is that there may be variation in the probability values of lower-level outcomes (that is, the probabilities over individual initial speeds), but the probability values of high-level outcomes (that is, probabilities over X_B and X_R)[25] are invariant under such variations. And the method of arbitrary functions is invoked to establish this claim. Indeed, the method shows that, given that the dynamics of a roulette wheel have very specific characteristics (they satisfy condi-

[25] The reader who prefers the probabilities to be defined literally on colour outcomes (and not merely on the set of initial speeds that yields these colour outcomes) can read the claims in this section as if the method has been extended in the way described at the end of §2.

tion **(a)**), any probability function p_i over individual speeds (so long as it satisfies condition **(b)**) will ascribe roughly the same probability of $1/2$ to the higher-level red and black outcomes. Thus the method explains why probabilities are invariant: they are high-level robust, because of the wheel's dynamics. Many authors think furthermore that invariance is a kind of objectivity. Thus according to them, the dynamics of the wheel make the probabilities in colour outcomes objective. In the rest of this section, I will examine and reject this view.

The view that the probabilities' robustness can be explained by the wheel's dynamics is a popular one. It can be found prominently in Butterfield (2011), who writes that these high-level probabilities "are robust in a vivid sense: the whole point of the method of arbitrary functions is that they are invariant under a choice of a density function from a wide class" (p. 1087). He claims that this hints that they are "objectively correct" (p. 1084). In a similar vein, Rosenthal (2012) "views objective probability as a high-level phenomenon that arises in deterministic contexts which are structured in a particular way" (p. 227). Abrams (2012) remarks that these "mechanistic" probabilities which are defined in terms of "causal structures" (p. 344) are "robust in the sense that it's difficult to alter [probabilities in lower-level outcomes] in such a way as to alter [probabilities in higher-level outcomes] to a significant degree" (p. 370). Finally, Strevens (2011) takes the method of arbitrary functions to help with "finding a basis for an outcome's probability in the properties of the physical dynamics that produce it." (p. 339).

There is clearly something true in the vicinity of what these authors claim. But, there are also two falsehoods, each either stated or implied by at least one author. These falsehoods are: (1) that robustness is a kind of objectivity; and (2) that the dynamics of the wheel play any role in establishing robustness. In the rest of this section, I expand on and refute each of these claims. It will follow that the method of arbitrary functions cannot help establish that the dynamics of a system objectivise probabilities, in this third sense.

Let me begin with the first claim, (1). The method of arbitrary functions can be used to show that the higher-level properties of a system do not depend sensitively

on its lower-level properties: the probabilities of colour outcomes do not depend sensitively on the probabilities of individual initial speeds. Thus the probabilities of colour outcomes may be called high-level *robust*. But this kind of robustness is not a form of objectivity. To see this, contrast our roulette wheel with what is known as a *chaotic* system, a system such that a small variation in initial conditions leads to a large variation in outcome state. By definition, many of the chaotic system's outcome properties are not robust: they depend very sensitively on the system's initial properties. But, I think it should be clear that these properties are in no way less objective than those of the roulette wheel. They may be less stable, and more difficult to determine, but they are not less objective. It follows that claim (1) is false.

Why then might someone have believed that robustness is a form of objectivity? The strongest reason I can think of is this. Robustness is a kind of invariance: the higher-level properties are invariant under changes in lower-level properties, if assumptions **(a)** and **(b)** are respected. Thus high-level robustness might also be called *lower-level invariance*. And furthermore, at least some kinds of invariance are forms of objectivity; such as (for example) perspective-invariance: something is objective in that sense if it is the same from all perspectives; and subjective otherwise.[26] Unfortunately, I have shown that the kind of invariance that the method of arbitrary functions is widely accepted to establish, lower-level invariance, is not a kind of objectivity. But the observation that robustness is a kind of invariance can nonetheless play a dual role: on the one hand it can help explain why (1) is held despite being false; and on the other hand it can suggest a potential strategy for arguing that the method can help with establishing objectivity, namely to identify a kind of invariance that the method does establish, besides lower-level invariance, and which does constitute a kind of objectivity.

Before concluding on (1), I want to explore and reject one way in which this strategy might be pursued. The hypothetical frequency interpretation, among other interpretations, suffers from what is known as the *reference class problem*.[27] In a

[26] This view is championed by Nozick (1998).

[27] There is a vast literature on this problem. The classic discussion begins with Reichen-

nutshell, the problem is this. The probability that a particular patient will get lung cancer depends on whether we consider him as an element of the set of all patients, or of the set of all smoker patients, or of the set of all heavy smoker patients, and so on. Thus what the probability values are depends on the reference class against which we consider this patient. As such, the probability that he will get cancer is reference-class-sensitive, and not reference-class-invariant. Furthermore, it is plausible to think that reference-class invariance is a kind of objectivity: the mention of the reference class is a mention of the modeller's particular perspective on the patient.[28] Could the method of arbitrary functions help with the reference class problem, and thus establish objectivity in the associated sense? We would have to interpret the function p_i as giving the probability of initial speeds given a particular reference class i. Furthermore, it would have to be the case that, no matter the reference class i, the shape of p_i respected assumption **(b)**, that is, did not vary too quickly. If these two constraints were satisfied, it would follow that the probability of red and black outcomes are roughly equal, no matter the reference class; and thus that the colour-outcome probabilities are objective in the associated sense. But this cannot be the case. For suppose that p_1 gives the probability for each initial speed that the wheel will be spun at that speed, and suppose furthermore that the chosen reference class is that which contains only this particular wheel spin. Then, p_1 will ascribe probability 1 to the speed at which the wheel is in fact spun, and probability 0 to all the other speeds, thus violating assumption **(b)**. It is therefore not true (and this may seem completely evident) that the probabilities of colour outcomes are roughly equal, for all reference classes.

Thus we arrive at claim (2), according to which the system's dynamics play a crucial part in establishing the robustness or level invariance of the colour probab-

bach (1949), whose PhD was, according to Butterfield (2011), on the method of arbitrary functions! For an insightful contemporary treatment, see Hájek (2007).

[28] This is suggested for instance by La Caze (2016), who writes that there is a sense in which hypothetical frequencies are not objective, namely, they "are not divorced from considerations of personal factors" (p. 358), in part because of the reference class problem.

ilities. My argument against this claim is a generalisation of an argument I already presented in §3. When applying the method of arbitrary functions to our roulette wheel, we interpreted the functions f and h as giving the dynamics or mechanics of the wheel. But the way in which the function is interpreted cannot play a role in determining what the method establishes; what plays a role is the more general point that f and g are coarse-graining functions. This is because, as I insisted in §2, the method is a mathematical theorem which holds independently of how its components are interpreted. Whether this function is interpreted as describing the mechanics of a system, or as describing the system's structure, is irrelevant. To see this more vividly, it suffices to notice that the method of arbitrary functions can be applied just as well to systems which are not dynamical, such as the urn example. In that case, the function from individual numbers to colours describes the (static) structure of the system of study. It follows that, contrary to popular opinion, whatever the method of arbitrary functions establishes, it does so independently of any system's dynamics (*qua* dynamics), or of any causal process.

Let me conclude on this section. The hope was that the method of arbitrary functions could help us show that the dynamics of a system can play a role in showing the system's probabilities to be objective, in a third sense of objective. But, in this section, that hope has been extinguished. Indeed, I have shown that the method of arbitrary functions can help establish that the probabilities in colour outcomes are robust. But I have also shown that robustness is not a form of objectivity, and that dynamics play no role in establishing robustness. (The fact that robustness is not a form of objectivity has dialectical implications for my paper. I claimed in §0 that I would show that the dynamics do not objectivise probabilities in any of the three senses of objectivity put forward by authors in the literature. For the first two senses, the properties discussed were indeed senses of objectivity, but the method did not show that dynamics gave rise to these properties. But for this third sense, it is for a deeper reason: because this sense of objectivity is not one after all—there is no objectivity in this sense.)

5.

Let me begin this paper's conclusion with a summary of its negative claims. I differentiated between three ways in which a probability function might be thought to be objective. Firstly, it may be ontically interpreted; secondly, it may be objectively evaluable (if it is interpreted as a credence function); thirdly, it may be high-level robust. The first two are indeed different senses in which a probability function may be objective. In this paper, I showed that the method of arbitrary functions cannot help show that probabilities can be objective in the first sense, and that it cannot establish that dynamics *qua* dynamics give rise to objectivity in the second sense. Then, I argued that, although the method of arbitrary functions does show that high-level probabilities are robust, this is not a sense of objectivity, nor are dynamics involved in establishing this fact. So there was a widespread hope, in the philosophical literature on probabilities, that the method of arbitrary functions could help show that probabilities in the sciences are objective in virtue of the system's dynamics, in one of three senses of objective. In this paper, I showed that this hope must be extinguished: the method can do no such thing.

I should note that it is consistent with everything I wrote in this paper for there to be another putative sense of objective, beyond those discussed in this paper, such that the method of arbitrary functions could be used to show that dynamics can objectivise probabilities, in this new sense of objective. I cannot think of any such new sense however, nor have I encountered such a sense in the literature. But, for those authors who remain committed to the thought that it is in the wheel's dynamics that the objectivity of probabilities lies, this option remains open in principle. Such an author would have to identify and give an account of this new sense, and show that it emerges from systems' dynamics.

My dialectical strategy has been to identify senses of objectivity in the work of authors in the philosophy of probability. But it should be clear, and I have hinted at, why the three genuine senses of objectivity mentioned in this paper can be deemed important for science, thereby explaining the interest that philosophers of probability have in them. One might want probabilities to be ontically interpreted,

on the grounds that science is about the world and not about what agents believe about the world. One might want them to be objectively evaluable, on the grounds that we cannot agree to disagree in science: there is instead a correct thing to believe. One might want them to be perspective-invariant (as briefly discussed in §4), on the grounds that the facts of science are true irrespective of the scientist's perspective on these facts. This suggests another task for the potential proponent of the new, dynamics-emergent sense of objectivity: to explain why this sense matters for science.

But even if there is no such new sense, it is untrue that the method is philosophically uninteresting. Its interest lies in the fact that it shows that particular kinds of dynamics are coarse-graining maps which yield robust probabilities over outcomes. In Butterfield's words, "the basic idea of the method of arbitrary functions [is] that intricate partitions of a sample space can wash out the peaks and troughs of an unknown density function, and secure robust probabilities" (p. 1090). Thus, the method shows that particular dynamics can be instances of coarse-graining maps that induce high-level robustness. The result is interesting, because conditions **(a)** and **(b)** are general enough that they apply to most actual roulette wheel spins. And Strevens (2011) claims, it can be applied beyond roulette wheel spins: a number of physical phenomena, such as coin flips, dice throws, and statistical mechanical phenomena might also exhibit these properties (though of course whether they do is an empirical matter). Thus, the method of arbitrary functions' value lies there: in showing that some robustness-yielding coarse-graining maps represent systems' dynamics. This value has been identified by the authors quoted in §4. But we must be careful not to take the infelicitous leap from this true claim to a false one about objectivity, or to another false one about the specialness of dynamics. This is what I have shown in this paper.

I will end this paper by making two remarks about what the method of arbitrary functions can be used to do, derivatively. I argued in §3 that the method cannot constrain credences directly, in providing a new norm of objective evaluability. However, it can play an indirect role, in the following way. As just explained,

the method of arbitrary functions can be used to establish that the frequencies of colour outcomes in roulette wheel spins are uniformly distributed, given some assumptions about the frequencies of individual initial speeds and about the structure of the system. Now suppose, as was first advocated by van Fraassen (1983), that an agent ought to align her credences in outcomes to their frequencies. Then, it would follow that an agent's credences over colour outcomes ought to be uniformly distributed. Thus the method can provide an indirect constraint on credences: by telling us that the frequencies in colour outcomes are what they are, it tells us what our credences in these outcomes should be.

Finally, I want to point to a last indirect possible use of the method of arbitrary functions, which as far as I am aware, has never been discussed. It is widely accepted that the results of a scientific experiment are only valuable insofar as they can be recovered; that is, it is possible to repeat the experiment and get the same result. But, in many cases, replicating the exact initial conditions of an experiment is difficult. The method of arbitrary functions indicates which probabilistic experiments are such that we should not worry too much about replicating the exact initial conditions. It tells us that if the target phenomenon satisfies conditions **(a)** and **(b)**, the specific initial conditions do not matter in a large enough number of trials: the outcome probabilities remain constant. In other words, the method helps us establish that and when the result of a probabilistic experiment is robust. But again, it would be an error to mistake this practical usefulness for conceptual significance.

Bibliography

Abrams, M. (2012). Mechanistic Probability. *Synthese 187*(2), 343–375.

Albert, D. Z. (2000). *Time and Chance*. Harvard University Press.

Berker, S. (2013a). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review 122*(3), 337–393.

Berker, S. (2013b). The Rejection of Epistemic Consequentialism. *Philosophical Issues 23*(1), 363–387.

Bernoulli, D. (1738/1954). Exposition of a New Theory on the Measurement of Risk. *Econometrica 22*(1), 23–36.

Bernstein, A. R. and F. Wattenberg (1969). Nonstandard Measure Theory. In W. A. J. Luxemburg (Ed.), *Applications of Model Theory to Algebra, Analysis and Probability*. Holt, Rinehard and Winston.

Bolker, E. (1966). Functions Resembling Quotients of Measures. *Transactions of the American Mathematical Society 124*(2), 292–312.

Bradley, R. (2017a). *Decision Theory with a Human Face*. Cambridge University Press.

Bradley, S. (2017b). Are Objective Chances Compatible with Determinism? *Philosophy Compass 12*(8).

Bradley, S. (2019). Imprecise Probabilities. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

Briggs, R. A. and R. Pettigrew (2020). An Accuracy-Dominance Argument for Conditionalization. *Noûs 54*(1), 162–181.

Bronfman, A. (2014). Conditionalization and Not Knowing That One Knows. *Erkenntnis 79*(4), 871–892.

Buchak, L. (2013). *Risk and Rationality*. Oxford University Press.

Butterfield, J. (2011). Less is Different: Emergence and Reduction Reconciled. *Foundations of Physics 41*(6), 1065–1135.

Canson, C. (forthcoming). Objectivity and the Method of Arbitrary Functions. *British Journal for the Philosophy of Science*.

Carnap, R. (1950). *Logical Foundations of Probability*. University of Chicago Press.

Carr, J. (2015). Epistemic Expansions. *Res Philosophica 92*(2), 217–236.

Carr, J. (2017). Epistemic Utility Theory and the Aim of Belief. *Philosophy and Phenomenological Research 95*(3), 511–534.

Carroll, L. (1895). What the Tortoise Said to Achilles. *Mind 4*(14), 278–280.

Chalmers, D. J. (2011). Frege's Puzzle and the Objects of Credence. *Mind 120*(479), 587–635.

Conee, E. and R. Feldman (2004). *Evidentialism: Essays in Epistemology*. Oxford University Press.

Davidson, D., J. C. C. McKinsey, and P. Suppes (1955). Outlines of a Formal Theory of Value. *Philosophy of Science 22*(2), 140–160.

de Bona, G. and J. Staffel (2018). Why Be Coherent? *Analysis 78*(3), 405–415.

de Finetti, B. (1937). La Prévision: Ses Lois Logiques, Ses Sources Subjectives. *Annales de l'Institut Henri Poincaré 7*, 1–68.

de Finetti, B. (1962). Does it Make Sense to Speak of Good Probability Appraisers? In I. J. Good (Ed.), *The Scientist Speculates. An Anthology of Partly-Baked Ideas*, pp. 357–64. Basic Books.

Descartes, R. (1641/1993). *Meditations on First Philosophy*. Hackett Publishing.

Dietrich, F., C. List, and R. Bradley (2016). Belief Revision Generalized: A Joint Characterization of Bayes's and Jeffrey's Rules. *Journal of Economic Theory 162*, 352–371.

Dorst, K. (forthcoming). Evidence: A Guide for the Uncertain. *Philosophy and Phenomenological Research*.

Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. MIT Press.

Easwaran, K. (2013). Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection. *Philosophy of Science 80*(1), 119–142.

Easwaran, K. (2014). Regularity and Hyperreal Credences. *Philosophical Review 123*(1), 1–41.

Easwaran, K. (2017). The Tripartite Role of Belief: Evidence, Truth, and Action. *Res Philosophica 94*(2), 189–207.

Easwaran, K. and B. Fitelson (2012). An 'Evidentialist' Worry About Joyce's Argument for Probabilism. *Dialectica 66*(3), 425–433.

Elga, A. (2007). Reflection and Disagreement. *Noûs 41*(3), 478–502.

Ellmann, L. (2019). *Ducks, Newburyport.* Galley Beggar Press.

Engel, E. (1992). *A Road to Randomness in Physical Systems.* Springer.

Feldman, R. and E. Conee (1985). Evidentialism. *Philosophical Studies 48*(1), 15–34.

Frege, G. (1892). On Sense and Reference. In P. Geach and M. Black (Eds.), *Translations from the Philosophical Writings of Gottlob Frege.* Oxford University Press. 1980.

Frigg, R. (2016). Chance and Determinism. In A. Hájek and C. Hitchcock (Eds.), *The Oxford Handbook of Probability and Philosophy.* Oxford University Press.

Galavotti, M. C. (2011). The Modern Epistemic Interpretations of Probability: Logicism and Subjectivism. In S. H. Dov M. Gabbay and J. Woods (Eds.), *Handbook of the History of Logic*, Volume 10, pp. 153–203.

Galavotti, M. C. (2016). The Origins of Probabilistic Epistemology. In A. Hájek and C. Hitchcock (Eds.), *The Oxford Handbook of Probability and Philosophy.* Oxford University Press.

Galavotti, M. C. (2018). Who is Afraid of Subjective Probability? In A. Giordani and C. de Florio (Eds.), *From Arithmetic to Metaphysics: A Path Through Philosophical Logic*, pp. 151–158. De Gruyter.

Gallow, J. D. (2017). Diachronic Dutch Books and Evidential Import. *Philosophy and Phenomenological Research 99*(1), 49–80.

Gallow, J. D. (forthcominga). A Subjectivist's Guide to Deterministic Chance. *Synthese.*

Gallow, J. D. (forthcomingb). Updating for Externalists. *Noûs.*

Glymour, C. (1980). *Theory and Evidence.* Princeton University Press.

Goodman, N. (1983). *Fact, Fiction, and Forecast.* Harvard University Press.

Greaves, H. (2013). Epistemic Decision Theory. *Mind 122*(488), 915–952.

Greaves, H. and D. Wallace (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind 115*(459), 607–632.

Haack, S. (1976). The Justification of Deduction. *Mind 85*(337), 112–119.

Hacking, I. (1967). Slightly More Realistic Personal Probability. *Philosophy of Science 34*(4), 311–325.

Hájek, A. (2003). What Conditional Probability Could Not Be. *Synthese 137*(3), 273–323.

Hájek, A. (2007). The Reference Class Problem Is Your Problem Too. *Synthese 156*(3), 563–585.

Hájek, A. (2012). Is Strict Coherence Coherent? *Dialectica 66*(3), 411–424.

Hájek, A. (2019). Interpretations of Probability. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

Hájek, A. (ms). Staying Regular.

Hájek, A. and H. Lin (2017). A Tale of Two Epistemologies? *Res Philosophica 94*(2), 207–232.

Hedden, B. (2015). Time-Slice Rationality. *Mind 124*(494), 449–491.

Henderson, L., N. D. Goodman, J. B. Tenenbaum, and J. F. Woodward (2010). The Structure and Dynamics of Scientific Theories: A Hierarchical Bayesian Perspective. *Philosophy of Science 77*(2), 172–200.

Hill, B. (2010). Awareness Dynamics. *Journal of Philosophical Logic 39*(2), 113–137.

Howson, C. (2000). *Hume's Problem: Induction and the Justification of Belief.* Oxford University Press.

Huber, F. (2017). On the Justification of Deduction and Induction. *European Journal for Philosophy of Science 7*(3), 507–534.

Hume, D. (1739/2000). *A Treatise of Human Nature.* Oxford University Press.

Hume, D. (1748/1999). *An Enquiry Concerning Human Understanding.* Oxford University Press.

Jackson, F. (1975). Grue. *Journal of Philosophy 72*(5), 113–131.

Jeffrey, R. C. (1965). *The Logic of Decision.* University of Chicago Press.

Jeffrey, R. C. (1968/1992a). Probable Knowledge. In *Probability and the Art of Judgment.* Cambridge University Press.

Jeffrey, R. C. (1991/1992b). Radical Probabilism. In *Probability and the Art of Judgment.* Cambridge University Press.

Johnson, D. (1991). Induction and Modality. *The Philosophical Review 100*(3), 399–430.

Joyce, J. (1998). A Non-Pragmatic Vindication of Probabilism. *Philosophy of Science 65*(4), 575–603.

Joyce, J. (1999). *The Foundations of Causal Decision Theory.* Cambridge University Press.

Joyce, J. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives 19*(1), 153–178.

Joyce, J. (2011). The Development of Subjective Bayesianism. In D. M. Gabbay, S. Hartmann, and J. Woods (Eds.), *Handbook of the History of Logic*, Volume 10.

Karni, E. and M.-L. Vierø (2013). Reverse Bayesianism: A Choice-Based Theory of Growing Awareness. *American Economic Review 103*(7), 2790–2810.

Karni, E. and M.-L. Vierø (2015). Probabilistic Sophistication and Reverse Bayesianism. *Journal of Risk and Uncertainty 50*(3), 189–208.

Kelly, T. (2003). Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research 66*(3), 612–640.

Keynes, J. M. (1921). *A Treatise on Probability.* Wildside Press.

Konek, J. and B. A. Levinstein (2019). The Foundations of Epistemic Decision Theory. *Mind 128*(509), 69–107.

La Caze, A. (2016). Frequentism. In C. Hitchcock and A. Hájek (Eds.), *The Oxford Handbook of Probability and Philosophy*, pp. 341–359. Oxford University Press.

Lange, M. (1999). Calibration and the Epistemological Role of Bayesian Conditionalization. *Journal of Philosophy 69*(9), 294–324.

Lange, M. (2002). Okasha on Inductive Scepticism. *Philosophical Quarterly 52*(207), 226–232.

Lange, M. (2011). Hume and the Problem of Induction. In D. M. Gabbay, S. Hartmann, and J. Woods (Eds.), *Handbook of the History of Logic*, Volume 10.

Leitgeb, H. and R. Pettigrew (2010a). An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of Science 77*(2), 201–235.

Leitgeb, H. and R. Pettigrew (2010b). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science 77*(2), 236–272.

Levi, I. (1985). Imprecision and Indeterminacy in Probability Judgment. *Philosophy of Science 52*(3), 390–409.

Lewis, D. K. (1980). A Subjectivist's Guide to Objective Chance. In R. C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability*, pp. 83–132. University of California Press.

Lewis, D. K. (1999). Why Conditionalize? In *Papers in Metaphysics and Epistemology*, pp. 403–407. Cambridge University Press.

Lyons, J. (2017). Epistemological Problems of Perception. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

MacAskill, W. and T. Ord (2020). Why Maximize Expected Choice-Worthiness? *Noûs 54*(2), 327–353.

Maher, P. (1995). Probabilities for New Theories. *Philosophical Studies 77*(1), 103–115.

Mahtani, A. (2017). The *Ex Ante* Pareto Principle. *Journal of Philosophy 114*(6), 303–323.

Mahtani, A. (forthcoming). Awareness Growth and Dispositional Attitudes. *Synthese*.

Millican, P. (1994). Hume's Argument Concerning Induction: Structure and Interpretation. In S. Tweyman (Ed.), *David Hume: Critical Assessments*. Routledge.

Myrvold, W. C. (2012). Deterministic Laws and Epistemic Chances. In Y. Ben-Menahem and M. Hemmo (Eds.), *Probability in Physics*, pp. 73–85. Springer.

Myrvold, W. C. (2016). Probabilities in Statistical Mechanics. In C. Hitchcock and A. Hájek (Eds.), *The Oxford Handbook of Probability and Philosophy*, pp. 573–600. Oxford University Press.

Norton, J. (2011). Challenges to Bayesian Confirmation Theory. In D. M. Gabbay, P. Thagard, and J. Woods (Eds.), *Handbook of the Philosophy of Science*, Volume 7.

Nozick, R. (1998). Invariance and Objectivity. *Proceedings and Adresses of the APA 72*(2), 21–48.

Okasha, S. (2001). What Did Hume Really Show About Induction? *Philosophical Quarterly 51*(204), 307–327.

Okasha, S. (2003). Probabilistic Induction and Hume's Problem. *Philosophical Quarterly 53*(212), 419–424.

Okasha, S. (2005). Bayesianism and the Traditional Problem of Induction. *Croatian Journal of Philosophy 5*(2), 181–194.

Paul, L. A. (2014). *Transformative Experience.* Oxford University Press.

Pettigrew, R. (2012). Accuracy, Chance, and the Principal Principle. *Philosophical Review 121*(2), 241–275.

Pettigrew, R. (2013a). A New Epistemic Utility Argument for the Principal Principle. *Episteme 10*(1), 19–35.

Pettigrew, R. (2013b). Accuracy and Evidence. *Dialectica 67*(4), 579–596.

Pettigrew, R. (2016a). *Accuracy and the Laws of Credence.* Oxford University Press.

Pettigrew, R. (2016b). Accuracy, Risk, and the Principle of Indifference. *Philosophy and Phenomenological Research 92*(1), 35–59.

Pettigrew, R. (2018). Making Things Right. In H. K. Ahlstrom-Vij and J. Dunn (Eds.), *Epistemic Consequentialism*, pp. 220–239. Oxford University Press.

Pettigrew, R. (2019). Epistemic Utility Arguments for Probabilism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy.*

Pettigrew, R. (2020). *Dutch Book Arguments.* Cambridge University Press.

Pettigrew, R. (forthcominga). Logical Ignorance and Logical Learning. *Synthese*.

Pettigrew, R. (forthcomingb). What is Conditionalization, and Why Should We Do It? *Philosophical Studies*.

Poincaré, H. (1896). *Calcul des Probabilités*. Carré et Naud.

Ramsey, F. P. (1926/1931). Truth and Probability. In *The Foundations of Mathematics*. Kegan Paul.

Reichenbach, H. (1949). *The Theory of Probability*. University of California Press.

Rosenthal, J. (2012). Probabilities as Ratios of Ranges in Initial-State Spaces. *Journal of Logic, Language and Information 21*, 217–236.

Roussos, J. (ms). Formal Philosophy as Modelling.

Russell, B. (2014). *A Priori* Justification and Knowledge. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

Salmon, W. (1990). Rationality and Objectivity in Science or Tom Kuhn Meets Tom Bayes. In C. W. Savage (Ed.), *Scientific Theories*, pp. 14–175. University of Minnesota Press.

Savage, L. (1973). Probability in Science: A Personalistic Account. In P. Suppes, L. Henkin, A. Joja, and G. C. Moisil (Eds.), *Logic, Methodology, and Philosophy of Science*, Volume 74, pp. 417–428.

Savage, L. J. (1954). *The Foundations of Statistics*. Wiley.

Schaffer, J. (2007). Deterministic Chance? *British Journal for the Philosophy of Science 58*(2), 113–140.

Schoenfield, M. (2015). Bridging Rationality and Accuracy. *Journal of Philosophy 112*(12), 633–675.

Schoenfield, M. (2017a). Conditionalization Does Not (in General) Maximize Expected Accuracy. *Mind 126*(504), 1155–1187.

Schoenfield, M. (2017b). The Accuracy and Rationality of Imprecise Credences. *Noûs 51*(4), 667–685.

Shimony, A. (1955). Coherence and the Axioms of Confirmation. *Journal of Symbolic Logic 1*(20), 1–28.

Shimony, A. (1970). Scientific Inference. In R. Colodny (Ed.), *The Nature and Function of Scientific Theories*, pp. 79–172. University of Pittsburgh Press.

Skyrms, B. (1980). *Causal Necessity.* Yale University Press.

Sliwa, P. and S. Horowitz (2015). Respecting *All* the Evidence. *Philosophical Studies 172*(11), 2835–2858.

Staffel, J. (2017). Should I Pretend I'm Perfect? *Res Philosophica 94*(2), 301–324.

Steele, K. and H. O. Stefánsson (ms.). Belief Revision for Growing Awareness.

Strevens, M. (2011). Probability Out Of Determinism. In C. Beisbart and S. Hartmann (Eds.), *Probabilities in Physics*, pp. 339–364. Oxford University Press.

Strevens, M. (2013). *Tychomancy.* Harvard University Press.

Stroud, B. (1977). *Hume.* Routledge.

Sylvan, K. (2018). Veritism Unswamped. *Mind 127*(506), 381–435.

Sylvan, K. (2020). An Epistemic Non-Consequentialism. *Philosophical Review 129*(1), 1–51.

Teller, P. (1973). Conditionalization and Observation. *Synthese 26*(2), 218–258.

van Fraassen, B. (1983). Calibration: A Frequency Justification for Personal Probability. In R. Cohen and L. Laudan (Eds.), *Physics, Philosophy, and Psychoanalysis.* Reidel.

van Fraassen, B. (1984). Belief and the Will. *Journal of Philosophy 81*(5), 235–256.

van Fraassen, B. (1989). *Laws and Symmetry.* Oxford University Press.

von Kries, J. (1886). *Die Principien der Wahrscheinlichkeitsrechnung.* Akademische Verlagsbuchhandlung von J. C. B. Mohr.

von Plato, J. (1983). The Method of Arbitrary Functions. *British Journal for the Philosophy of Science 34*(1), 37–47.

Wenmackers, S. (2019). Infinitesimal Probabilities. In R. Pettigrew and J. Weisberg (Eds.), *The Open Handbook of Formal Epistemology.* PhilPapers Foundation.

Wenmackers, S. and J.-W. Romeijn (2016). New Theory About Old Evidence. *Synthese 193*(4), 1225–1250.

Werndl, C. (2015). Determinism and Indeterminism. In P. Humphreys (Ed.), *The Oxford Handbook of Philosophy of Science*. Oxford University Press.

White, R. (2009). Evidential Symmetry and Mushy Credence. *Oxford Studies in Epistemology 3*(1), 161–186.

Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford University Press.

Yalcin, S. (2018). Belief as Question-Sensitive. *Philosophy and Phenomenological Research 97*(1), 23–47.

Zabell, S. (2016). Symmetry Arguments in Probability. In A. Hájek and C. Hitchcock (Eds.), *The Oxford Handbook of Probability and Philosophy*. Oxford University Press.