Global Constraints on Imprecise Credences: Solving Reflection Violations, Belief Inertia, and Other Puzzles

Sarah Moss
ssmoss@umich.edu

A lot of foundational work in formal epistemology proceeds under the assumption that subjects have precise credences. The traditional requirement of probabilistic coherence presupposes that you have precise credences, for instance, and it says that your precise credence function must satisfy the probability axioms. The traditional rule for updating says that when you get evidence, you should modify your precise credence function by conditionalizing it on the information that you learn. Meanwhile, advocates of imprecise credences challenge the assumption behind these rules. They argue that your partial beliefs are best represented not by a single function, but by a set of functions, or *representor*.[2] The move to imprecise credences leaves many traditional requirements of rationality surprisingly intact, as fans of imprecise credences often simply reinterpret these rules as applying to the individual functions in your representor. In order for your imprecise credences to be rational, each member of your representor must satisfy the probability axioms. In order for you to update rationally, your later representor must contain just those functions that result from conditionalizing each member of your representor on the information you learn.[3]

However, for agents with imprecise credences, the requirements of rationality needn't take this form. Whether you are rational might just as easily depend on *global*

---

1. I am grateful to Eric Swanson for several insightful comments that prompted the writing of this paper, and for many subsequent conversations about the central ideas in it. Thanks also to Jim Joyce, Teddy Seidenfeld, and an anonymous referee for helpful discussion.

2. The term 'representor' is from van Fraassen 1990. For early discussions of imprecise credence models, see Smith 1961, Levi 1974, and Williams 1976.

3. This updating rule is part of the definition of an *imprecise probability model* in the sense of Joyce 2010. For further discussion, see the literature cited in footnote 32.

features of your representor, features that can't be reduced to each member of your representor having a certain property. Global features of your representor are like the properties attributed by collective readings of predicates such as 'lift the piano'. What it takes for a group of people to lift a piano is not the same as what it takes for each individual member of the group to lift it. Similarly, what it takes for an imprecise agent to be rational might not be for each member of her representor to satisfy familiar constraints on precise credence functions. To take the point further, imagine a band director commanding a marching band to spread out to fill a football field. This command is global in an especially strong sense—namely, no individual could possibly satisfy it. Similarly, for fans of imprecise credences, the requirements of rationality could in principle include rules that no precise agent could possibly satisfy.

This paper is an extended investigation of global rules of rationality. Some rules surveyed in this paper are rules analogous to the command to lift a piano, and some are analogous to the command to spread out to fill a football field. In section 1, I state formal definitions for both of these kinds of global constraints, and I address relevant questions about how to interpret the formalism of imprecise credences. In the remainder of the paper, I describe three applications of global constraints, using my ideas to solve problems faced by fans of imprecise credences. Section 2 discusses cases in which it seems like imprecise agents are forced to make bad choices about whether to gather evidence. Section 3 discusses the problem of belief inertia, according to which certain imprecise agents are unable to engage in inductive learning.[4] Finally, section 4 answers the objection that many imprecise agents are doomed to violate the rational principle of Reflection.[5] These three applications are modular enough that readers interested in one particular problem may skip to my discussion of it after reading section 1.

A note of clarification: in discussing global requirements of rationality, I am playing a defensive game on behalf of fans of imprecise credences. I am not aiming to prove that imprecise credences are sometimes rationally required, or even that they are rationally permissible. Rather, I am aiming to demonstrate that fans of imprecise credences have more argumentative resources at their disposal than previously thought, resources brought out by the observation that the rules of rationality could be global in character. Imprecise credence models can support a much broader range of rational requirements than precise credence models, and fans of imprecise credences can benefit from understanding this flexibility and taking better advantage of it.

4. For an introduction to the problem of belief inertia, see §3.2 of Bradley & Steele 2014.
5. For a prominent statement of this objection, see White 2010.

1 Two notions of globalness

Let a *representor* be a set of probability measures, and let a *constraint* be a set of representors. We define the notion of a *pointwise* constraint as follows:

C is *pointwise* if and only if: there is some set of probability measures S such that for every representor R, $R \in C$ if and only if $R \subseteq S$.

When a constraint is pointwise, we can figure out whether it contains a representor just by testing to see whether every probability measure contained in that representor has a certain property. For example, say that you have .5 credence that a certain fair coin will land heads, although your credences in other propositions are less precise. Then your representor will be a member of a certain constraint, namely the set of representors whose members agree that it is .5 likely that the coin will land heads. This is a pointwise constraint, satisfied by your representor in virtue of the fact that every one of its members assigns .5 probability to the coin landing heads.

With this notion in hand, we can define our first notion of globalness:

C is *global* if and only if: C is not pointwise.

For example, consider the set of representors that contain *at least one* probability measure that assigns .5 to the coin landing heads. This constraint does not correspond to a test on the individual members of a representor, as evidenced by the fact that a representor can satisfy this constraint, while a proper subset of that same representor fails to satisfy it. Pointwise constraints are like distributive readings of predicates, which are satisfied in virtue of every member of a group having a certain property, while global constraints are more like collective readings.[6]

In order to spell out another useful characterization of this first notion of globalness, we must make a small detour and say more about how to interpret the formalism of representors and constraints. What exactly does your representor represent? This question is best answered by analogy. According to a traditional model of full beliefs, you believe a proposition just in case it contains every world that is doxastically possible for you.[7] Analogously, your representor can be used to model your *probabilistic beliefs*—that is, your credences, conditional credences, comparative probability judgments, and so on. We can think of these probabilistic beliefs as attitudes toward sets of probability spaces, or *probabilistic contents*.[8] For instance, you have .6 credence that

6. For an overview of the distributive–collective distinction, see Champollion 2020, §2–3.
7. Classic discussions of this model of belief include Hintikka 1962, Stalnaker 1984, and Lewis 1986a.
8. For an extended defense of the claim that probabilistic beliefs are attitudes toward probabilistic contents, see §1.2–3 and §3.6 of Moss 2018.

Jones smokes in virtue of standing in the belief relation to a certain set of probability spaces, namely those that assign .6 probability to the proposition that Jones smokes. Just as you believe a proposition if and only if it contains every one of your doxastic possibilities, you believe a probabilistic content if and only if it contains every member of your representor. For instance, you have .6 credence that Jones smokes if and only if every member of your representor assigns .6 probability to the proposition that Jones smokes.[9]

Just as a set of worlds can represent your full beliefs, and your representor can represent your probabilistic beliefs, these same models can also represent another doxastic attitude. Given any content such that you could believe it or believe its complement, there is also a third attitude that you can hold toward that content—namely, the attitude of suspending judgment. As Friedman 2013 convincingly argues, suspending judgment is a genuine attitude, not the mere absence of belief or disbelief. As Friedman puts it, suspending judgment about a content is an attitude that "expresses or represents or just is [your] neutrality or indecision" about that content (180). According to traditional models of full belief, you suspend judgment about a proposition just in case it contains some but not all of your doxastic possibilities. According to imprecise credence models, agents with imprecise credences suspend judgment about probabilistic contents.[10] Levi 1980 describes this interpretation of imprecise credence models by saying, "[c]redal ignorance entails suspension of judgment between alternative systems of evaluations of hypotheses with respect to credal probability" (185). As Kaplan 2010 puts it, imprecise credence models represent "a doxastic option, indecision, whose cogency orthodox Bayesian Probabilism wrongly refuses ever to countenance" (49). To be more specific, you suspend judgment about a probabilistic content just in case it contains some but not all of the probability measures in your representor. For example, suppose you are wondering whether it is at least .6 likely that Jones smokes. If you do not know whether it is at least .6 likely that she smokes, you may want to avoid taking a stand on this probabilistic question. In order to suspend judgment about it, you must have some representor members that assign at least .6 probability to Jones smoking, as well as some representor members that do not. Imprecision *just is* the suspension of probabilistic judgment.

Having spelled out this interpretation of imprecise credence models, we can iden-

9. In the rest of this paper, for sake of simplicity, I will talk about sets of probability *measures* rather than sets of probability *spaces*. Although probabilistic contents are defined to be the latter rather than the former in Moss 2018, this difference does not matter for my arguments. Also, I set aside potential differences between what is represented by a probability measure and by its singleton set, assuming that the belief states of precise agents may be represented by either object.

10. Although this is my preferred interpretation of imprecise credence models, there are notable alternatives which I am setting aside for present purposes. For discussion, see §2.10.3 of Walley 1991.

tify a second useful characterization of the notion of a pointwise constraint defined above. A pointwise constraint contains the representors of all and only those agents who *believe a certain probabilistic content*. By the definition given above, a constraint is pointwise just in case it is the power set of some set S of probability measures. Being a member of S is the test that each member of your representor must pass in order for your representor to satisfy the constraint. When each member of your representor is indeed contained in the set S, that just amounts to saying that S is a probabilistic content that you believe.

In addition to identifying certain constraints as global, we can also identify a special subset of global constraints, namely those that do not contain the representors of any precise agents.[11] These constraints are global in an especially strong sense:

C is *strongly global* if and only if: for all $R \in C$, $|R| > 1$

All other global constraints are merely *weakly global*:

C is *weakly global* if and only if: C is global, and for some $R \in C$, $|R| = 1$

Some properties of groups—filling a football field, for instance—are properties that only groups can have. Strongly global constraints are like these properties. By contrast, weakly global constraints are like the property of lifting a piano. Although groups can lift pianos, so can very strong individuals. But the property of lifting a piano is still global in an interesting sense—namely, because a group does not have this property in virtue of each member of the group having it.

A word of caution: although it can be helpful to compare global constraints and collective readings of predicates, it is important not to overstate the analogy between them. For instance, one might at first be tempted to assume that any global constraint must contain at least one non-singleton set. But this assumption is false. For instance, consider the set of representors containing exactly one probability measure—that is, the constraint corresponding to the rational requirement to have precise credences. This set is a global constraint. Having exactly one measure in your representor does not involve each of your representor members having a certain property. The notion of a global constraint essentially depends on the richness of imprecise credence models, but not because every global constraint must contain the representor of some imprecise agent.

11. Any non-empty constraint of this sort must be global, since any non-empty pointwise constraint will contain some precise credence function(s)—namely, the representor(s) of each precise agent who believes the corresponding probabilistic content.

2 The problem of cheap evidence

In sections 3 and 4, I use the foregoing notions of globalness to solve well-known problems for imprecise credences. But first, I want to introduce and address a serious problem that has not been widely discussed in the literature. The problem is that without global constraints on rationality, imprecise agents will be forced to make questionable decisions about whether to gather evidence.

2.1 *Phone a Friend*

Consider the following example:

> *Phone a Friend*: You are about to be offered a chance to guess whether p is true. If your guess is correct, you will win 100 dollars. If your guess is incorrect, you will lose 100 dollars. Also, before you face this offer, you have the option of paying 20 dollars right now to phone a friend and find out whether p is true.

Suppose that you are certain that if you phone a friend and find out whether p is true, you will guess accordingly. What if you decide not to phone a friend? Let *Later p* be the material conditional proposition that if you remain uninformed, you will later guess that p is true. Let *Later not-p* be the proposition that if you remain uninformed, you will guess that p is not true. Suppose that your representor contains just two probability measures, m_1 and m_2, and that these measures have the following features:

$$m_1(p) = .99,\ m_1(\textit{Later } p) = 1,\ m_1(\textit{Later not-}p) = 0$$

$$m_2(p) = .01,\ m_2(\textit{Later } p) = 0,\ m_2(\textit{Later not-}p) = 1$$

The members of your representor strongly disagree about the likelihood of p and also about the likelihood that you will later guess that p is true. Their opinions about how you *will* act reflect their first-order credences that determine how you *should* act. This sort of dependence in your imprecise credences is described and defended by Williams 2014 in response to another diachronic decision puzzle. As Williams puts it, the idea is that "each credence assumes that the agent will do what is rational" by the lights of that credence function (26)—which in this case, means guessing that p is true just in case it is likely that p is true.

Unfortunately, *Phone a Friend* presents a problem for imprecise credence fans. According to each member of your representor, you should not pay 20 dollars to find out whether p is true. That's because according to each representor member, the expected utility of foregoing the phone call and making an uninformed guess is

98 dollars, whereas the expected utility of making an informed guess is merely 80 dollars. As a result, almost every decision theory for imprecise agents will say that it is *impermissible* for you to make an informed guess. As Joyce 2010 explains, there is a "consensus among proponents of the imprecise model" that the rules for rational decision making "should never recommend one act over another when every member of your committee says that the utility of the latter exceeds that of the former" (311). Since your representor members unanimously agree that you ought to forego the phone call, it follows that you ought to forego the phone call.

This is a counterintuitive result. Although the members of your representor agree that you should make an uninformed guess about whether p is true, they have wildly divergent opinions about the expected value of particular options that you will later face, such as the act of guessing that p is true. In light of this fact, it is intuitive to think that rationality should at least *permit* you to gather evidence that would settle this dispute and tell you what to guess in order to win 100 dollars. Call this the *problem of cheap evidence*. To sum up, it ought to be permissible for an imprecise agent to pay for evidence that will help her make up her mind about decisions, as long as the evidence is cheap enough to be worth it. But absent any relevant constraints on rational credences, it is unclear how to secure this result.[12]

2.2 Convexity to the rescue?

How should the fan of imprecise credences respond? There are a couple of options, both of which involve endorsing global requirements on imprecise agents, rules that forbid rational agents from having a representor containing only m_1 and m_2. The first requirement is a rule proposed by Levi 1980—namely, that your representor must be *convex* in the following sense:

R is *convex* if and only if: for all $f, g \in R$ and $0 \leq \lambda \leq 1$, $\lambda f + (1-\lambda)g \in \mathbb{R}$

A convex representor contains all linear averages of its members. As a result, the credences assigned by the members of a convex representor take a special form. Where R is a representor, and p is a proposition, let us define $R[p] =_{df} \{m(p) : m \in R\}$. Speaking loosely, $R[p]$ is "the imprecise credence assigned by R to p." If R is convex, then for any proposition p, $R[p]$ will be an interval of real numbers.

12. The problem of cheap evidence is distinct from the *problem of free evidence* discussed in the literature. The former problem is that it is intuitively permissible for agents to pay for certain evidence, whereas the latter problem is that it is intuitively obligatory for agents to gain evidence that they do not have to pay for. The former problem arises when each representor member is confident that the imprecise agent will make decisions that are correct from her perspective, whereas the problem of free evidence resolves itself under these same conditions, as discussed in §3.1 of Bradley & Steele 2016.

The set of convex representors is a global constraint. A representor can be convex while a proper subset of it is not, and so the constraint of convex representors is not the power set of any set of probability measures. Among the many global requirements that might govern imprecise agents, convexity is certainly one of the first that comes to mind.[13] Roughly speaking, the idea of this requirement is that your representor members must "fill in any gaps in the football field." In other words, if your representor includes some probability measures, it must also include all those between them.

Requiring agents to have convex credences forestalls our counterintuitive verdict about *Phone a Friend*, and indeed, it provides a general strategy for solving problems of this sort. As long as your representor is convex, any pair of confident representor members like m_1 and m_2 will be accompanied by a moderate probability measure that assigns .5 to the proposition p that you might be acting on later. According to this third probability measure, the later option to guess that p is worth much more if you first find out whether p is true, and so you should be willing to pay a lot to gain that information. As long as your representor contains moderate probability measures, you will not be rationally required to forego gathering evidence in cases like *Phone a Friend*.

Since the convexity requirement is endorsed by some imprecise credence fans, it is important to appreciate that it solves the problem of cheap evidence. However, the requirement itself is highly controversial. The most prominent argument against the requirement is due to Jeffrey 1987. According to Jeffrey, you can judge that propositions are irrelevant to each other without having settled opinions about either. For instance, you can be certain that whether there is water on Mars is independent of whether a certain coin landed heads, even if you have maximally imprecise credences about both propositions. But as Jeffrey points out, the convexity requirement precludes rational agents from having this particular combination of beliefs:

> The bare judgment of irrelevancy would be represented by the set $I = \{P|P(AB) = P(A)P(B)\}$, but by no proper subset, and by no one member. Levi disallows such judgments: he requires the sets that represent indeterminate probability judgments to be convex. (586)

The set of probability measures representing the bare judgment of irrelevancy is not convex, and so Levi says that it cannot constitute the representor of any rational agent.

As I see it, Jeffrey's observation alone does not yet constitute a strong argument against the convexity requirement. Fans of convexity could reasonably complain that

13. Convexity is championed by Levi 1980, ch. 9; it is also an assumption of the decision theory defended in Gilboa & Schmeidler 1989 and the maximality theorem proved in Walley 1991.

it is unrealistic to assume that a rational agent could believe a bare judgment of irrelevancy about some propositions while lacking any other belief whatsoever. But on behalf of opponents of convexity, I want to propose a way to strengthen Jeffrey's argument, challenging convexity without making this unrealistic assumption. Here is a much weaker assumption: a rational agent can believe that A and B are irrelevant to each other, while still having imprecise credences in A and B. For fans of imprecise credences, this certainly seems like an unobjectionable combination of belief states, and yet the convexity requirement rules out any such state as rationally impermissible. The set I defined by Jeffrey has the following significant feature: for any $R \subseteq I$, R is convex only if one of the following holds: all of its members assign the same credence to A, or all of them assign the same credence to B.[14] Hence the convexity requirement entails that any rational agent with a representor in I must have a precise credence in at least one of A or B.

To sum up, requiring representors to be convex would indeed solve the problem of cheap evidence. Levi and other proponents of convexity should welcome this motivation for their global constraint. But given our strengthened challenge for convexity, one might also reasonably doubt whether imprecise agents are rationally required to have convex representors.[15]

2.3 A global independence requirement

In light of the controversial nature of the convexity requirement, I want to state an additional solution to the problem of cheap evidence. Rather than faulting your credences in p for not being convex, we can instead fault them for being intimately connected with your credences about how you will later act. Imagine filling out the details of the *Phone a Friend* case by saying that there is no fact of the matter that settles exactly which member of your representor will determine how you respond in any given decision situation. If that is right, then there is no evidence that could settle whether you will guess that p if you do not phone a friend. Even from the perspective of an individual member of your representor, there is no more reason to suppose that you will later guess correctly than that you will guess incorrectly. Accordingly, rationality should require you to have imprecise credences about how you will guess—and moreover, it should require these credences to be *independent* of your first-order beliefs about p.

How should we understand this independence requirement? Crucially, the requirement does not involve making the sort of bare judgment of irrelevancy discussed

14. See claim 1 of the appendix.
15. For additional critical discussion of convexity, see Kyburg & Pittarelli 1996.

by Jeffrey 1987. In Jeffrey's sense, you do believe that p and *Later p* are irrelevant to each other in *Phone a Friend*. As you decide whether to phone your friend, each member of your representor assigns an extreme credence to *Later p*. Hence each member of your representor considers p and *Later p* to be trivially independent, so your representor is indeed contained in the following constraint:

$$I_p =_{df} \{R : R \subseteq \{m | m(p \wedge Later\ p) = m(p)m(Later\ p)\}\}$$

We cannot solve the problem of cheap evidence merely by requiring your representor to satisfy this pointwise constraint.

By contrast, a global interpretation of the relevant independence requirement yields a successful solution to our problem. Consider the following requirement:[16]

Global Independence For any probability measures m_1 and m_2 in your representor, there must be a third representor member m_3 such that $m_3(p) = m_1(p)$, and such that m_3 is certain that you will later act as prescribed by your currently having m_2 as your credence function.

When you satisfy *Global Independence*, it is as if your representor members themselves have imprecise credences about how you will act. Strictly speaking, of course, we cannot require the precise functions of your representor to suspend judgment about anything. But *Global Independence* imposes a functionally equivalent requirement—namely, that any precise credence function in your representor is accompanied by many counterpart functions, each of which is certain that you will act on the recommendations of a different representor member. The set of these counterpart functions is the representor of an agent who suspends judgment about how you will act. Rather than demanding that you believe that p and *Later p* are independent, *Global Independence* demands that you suspend judgment about each proposition in light of the other. This fact explains why *Global Independence* is a global requirement of rationality—namely, because it does not require probabilistic belief, but rather the suspension of probabilistic judgment.

By contrast with I_p, the constraint of *Global Independence* is not satisfied by your representor in *Phone a Friend*. We could produce a representor that satisfies this constraint by adding the following probability measures to your representor:

$$m_3(p) = .99,\ m_3(Later\ p) = 0,\ m_3(Later\ not\text{-}p) = 1$$

16. A number of global notions of independence could be used to solve our puzzle. See Moss 2015 for a precursor to the above independence requirement, as well as discussion of another puzzle that it could be used to solve. For further discussion of global notions of independence, see §3 of Couso et al. 2000 and §3.1 of Cozman 2012.

$m_4(p) = .01,\ m_4(\mathit{Later}\ p) = 1,\ m_4(\mathit{Later}\ \mathit{not}\text{-}p) = 0$

But as soon as these probability measures are added, the act of declining to gather evidence about p will no longer have maximal expected value according to every member of your representor. According to m_3 and m_4, the expected utility of making an uninformed guess is -98 dollars, whereas the expected utility of making an informed guess is 80 dollars, so there is no longer any simple argument for the claim that you are rationally required to forego gathering evidence.

Global Independence serves as a useful template for independence requirements on imprecise credences, applicable in cases where independence should be understood in terms of suspending judgment, rather than believing a probabilistic content that represents irrelevancy. *Phone a Friend* motivates one rational requirement of this sort—namely, the requirement that your first-order credences be independent of your credences about your future actions. Of course, your predictions about your actions may sometimes depend on your first-order beliefs—such as, for instance, when you are offered a bet at long odds on the proposition that you will at some point accept at least one bet at long odds. But in many normal cases, an imprecise agent will not have much evidence about which members of their representor will end up governing their actions. In the absence of relevant evidence, an agent will be rationally required to suspend judgment, and they will thereby avoid the problem of being required to forego sensible acts of evidence gathering.

3 The problem of belief inertia

Fans of imprecise credences sometimes suggest that rational agents can have *radically imprecise* credences in a proposition, i.e. credences that span the range from 0 to 1. At first glance, however, this appears to raise a problem. In certain circumstances, rational agents with radically imprecise credences will retain those credences, no matter what evidence they conditionalize on. In other words, some radically imprecise credences are *inert* with respect to relevant evidence.[17] As Walley 1991 observes, "[i]f the vacuous previsions are used to model prior beliefs about a statistical parameter for instance, they give rise to vacuous posterior previsions" (93). Rinard 2013 gives an example:

> [C]onsider an urn about which you know only the following: either all the marbles in the urn are green (H1), or exactly one tenth of the marbles are green (H2)... if

17. Formally, a representor R is *inert* with respect to a constraint C and a set of propositions E if and only if for all $p \in E$, $\{m|_p : m \in R\} \in C$, where $m|_p$ is the result of conditionalizing the measure m on the proposition p. The set E is often implicitly determined by context to be the set of evidence propositions that an agent might learn.

> your initial credence in H1 is $(0, 1)$, it will remain there. It will be impossible for you to become confident in H1, no matter how many marbles are sampled and found to be green. (160–1)

In this example, your imprecise credence in H1 will be inert with respect to the proposition that all of the sampled marbles have been green, no matter how many marbles have been sampled. In this sense, having radically imprecise credences can preclude inductive learning. And yet, it seems rationally impermissible for an agent to be incapable of learning from experience. As White 2010 puts the point, "Maximally mushy credences are immovable! This result is entirely unacceptable" (184). This apparent problem for fans of imprecise credences has come to be known as the *problem of belief inertia*.[18]

In response to the problem of belief inertia, Joyce 2010 proposes that rational imprecise credences must satisfy two constraints:

> Perhaps the right way to secure inductive learning is to sharpen your credal state by (a) throwing out all the pigheaded committee members... and (b) silencing "extremist" elements by insisting that each committee member assign a credence to [H1] that falls within some sharpened interval. (291)

The exact details of this proposal do not matter for our purposes. The point is that Joyce is attempting to derive a global requirement of rationality from pointwise requirements. The idea is that there are certain rational constraints on precise credence functions, namely constraints against pigheaded and extremist credences. Fans of imprecise credences should endorse these requirements as constraints on the individual members of your representor. According to Joyce, any representor satisfying these pointwise constraints will also satisfy the constraint of not being inert with respect to assigning radically imprecise credences to a proposition.

Unfortunately, Vallinder 2018 demonstrates that Joyce's proposal does not solve the problem of belief inertia. Vallinder produces a representor that satisfies the pointwise constraints that Joyce proposes, although it is still incapable of inductive learning. The crux of Vallinder's argument is that you can be stubborn in your credences without being maximally imprecise. Even in the absence of pigheaded and extremist representor members, your credences can be inert with respect to moderate imprecise credence assignments. Vallinder concludes that the problem of belief inertia remains, since "even this weaker form of belief inertia means that no matter how much evidence the agent receives, she cannot converge on the correct answer with any greater

18. This terminology is due to Bradley 2012, though the problem is widely discussed by earlier authors. For a classic discussion of belief inertia, see §13.2 of Levi 1980. For more recent statements of the problem, see Weatherson 2008, §4.2; Bradley 2012, §4.6.3; Rinard 2013, p. 160ff.; and Vallinder 2018.

precision than is already given in her prior credal state" (1216).

How, then, should imprecise credence fans solve the problem of belief inertia? As I see it, the correct response does not involve deriving anti-inertia requirements from pointwise requirements of rationality. Anti-inertia requirements on imprecise credences are indeed intimately connected with rational requirements on precise agents, but not because the former consist in the pointwise application of the latter. Rather, many anti-inertia requirements on imprecise agents are genuinely global requirements. They are not derived from rules for precise agents, because they are themselves the direct analogs of traditional rules against inert belief states.

To back up a step, note that in discussions of rational requirements governing full beliefs, it is often taken for granted that you should not be stubborn in your beliefs. As Quine 1951 puts it, "no statement is immune to revision" (40), not even obvious statements such as the law of excluded middle. This statement is generally interpreted as a normative claim. In order to be a rational agent, none of your beliefs should be so strong that you would retain it in the face of any counterevidence whatsoever. In short, you should be open-minded rather than stubborn.

When it comes to agents with precise credences, the most familiar rule against stubbornness is the rule of *Regularity*, which requires rational agents to assign positive credence to any epistemically possible proposition. Arguments for Regularity often take it for granted that rational agents should be willing to revise their beliefs. Here are some examples:

> Absolute certainty is tantamount to a firm resolve to never change your mind no matter what, and that is objectionable.[19]
>
> [An agent] who started out with an irregular credence function (and who then learned from experience by conditionalizing) would stubbornly refuse to believe some propositions no matter what the evidence in their favor.[20]
>
> The idea behind these arguments [for Regularity] is that being doxastically stubborn might make us miss out on good beliefs that we could have in the future, whereas being open-minded allows us to have those beliefs.[21]
>
> Moral: Keep the mind open, or at least ajar.[22]
>
> [R]egularity… is meant to capture a form of open-mindedness and responsiveness to evidence.[23]

19. Lewis 1981, p. 14.
20. Lewis 1980, p. 268.
21. Singer 2019, p. 291.
22. Edwards et al. 1963, p. 211.
23. Hájek 2019, §3.3.4.

To sum up, advocates of Regularity typically defend it by appealing to the intuitive thought that rational agents should be open-minded rather than stubborn.[24]

This intuitive thought supports other rational requirements on precise agents as well. For instance, a natural extension of Regularity would not only forbid you from being dogmatic about which propositions you believe with maximum confidence, but also forbid you from being dogmatic about other probabilistic beliefs. Also, it is intuitive to think that rational agents should not only be open to revising their beliefs, but also open to revising other attitudes, such as the attitude of suspending judgment about a question. Following Weatherson 2015, we can distinguish the following rational requirements:

Open-Minded	Any time an agent is confident in a proposition, there is some evidence she could get that would make her lose confidence in it.
Evidence-Responsive	For any proposition, there is some evidence the agent could get that would make her confident in it. (9)

Just as rational agents must be able to both gain and lose full beliefs, rational agents must be able to both gain and lose probabilistic beliefs, such as attitudes of confidence. Just as any thread might eventually be disentangled from your web of probabilistic belief, many external threads might eventually be woven into it.

The traditional rule of Regularity governs precise agents. However, this rule against stubbornness can be extended to imprecise agents, and the same goes for more general rules against stubbornness. For any given constraint on imprecise credences, there is a second constraint containing all and only those representors that are not inert with respect to the first constraint. Rules against stubbornness are constraints of this second sort, many of which are global constraints. For example, consider the constraint containing representors of agents that believe that it is more than .5 likely that a certain coin landed heads. Although this is a pointwise constraint, the set of representors that are not inert with respect to it is a global constraint. Although believing that a coin probably landed heads just amounts to each member of your representor assigning at least .5 probability to the proposition that it landed heads, *being such that you could stop believing this content* does not correspond to any pointwise test.[25]

24. I am not defending Regularity in this paper. It is controversial whether an agent must satisfy Regularity in order to avoid being stubborn; see Easwaran 2014 for an opposing view. The focus of my discussion is the underlying assumption that rational agents should avoid being stubborn, as it is this widely shared assumption that I want to extend to imprecise agents.

25. See claim 2 of the appendix.

Let us return to the problem of belief inertia. In his criticism of Joyce, Vallinder assumes that when it comes to imprecise agents, the impermissibility of belief inertia must be grounded in other rational requirements. But why should we accept this assumption? When it comes to precise agents, rules against stubbornness are not grounded in more basic requirements. The rule of Regularity is compelling because it follows from the intuitive idea that rational agents should be open-minded. Rules against inert credences—precise or imprecise—can be defended on just the same grounds. The impermissibility of inert credences is not derived from pointwise requirements, but from more general global principles. The phenomenon of belief inertia is an instance of the broader phenomenon of having stubborn imprecise beliefs, which are just as irrational as stubborn precise beliefs, and for just the same reasons. It is a general fact that rational agents are not incapable of learning, and it follows from this general fact that rational imprecise agents do not have inert credences that are maximally imprecise.

4 Violations of Reflection principles

Like many rules of rationality, the principle of Reflection has traditionally been interpreted as imposing a constraint on the credence functions of precise agents. As introduced by van Fraassen 1984, the principle states that your conditional credence in a proposition, conditional on your assigning it credence r at some later time, must equal this same real number r (244). In other words:

Precise Reflection $\quad Cr_0(p|Cr_1(p) = r) = r$

As many authors have noted, this principle stands in need of qualification.[26] *Precise Reflection* should not govern your credences when you believe that you might forget information, for instance, or when you fear that you might learn false information or fail to update rationally. Briggs 2009 makes the helpful observation that a suitably qualified version of *Precise Reflection* is simply a consequence of the Kolmogorov axioms, and so the former may be considered just as unobjectionable as the latter.

In the context of *Precise Reflection*, Cr_0 and Cr_1 are precise credence functions, mapping propositions to real numbers. How should this principle be extended to constrain representors, objects that are not even functions defined on propositions? At first glance, this question might appear to have an obvious answer—namely, that Reflection constrains your current imprecise conditional credences in a proposition,

26. See Christensen 1991, Talbott 1991, Maher 1992, and Bovens 1995.

given hypotheses about your later imprecise credence in it. In other words, one might be tempted to extend Reflection as follows:[27]

Value Reflection $\quad R_0[p|R_1[p] = S] = S$

In support of *Value Reflection*, White 2010 says, "It is natural to suppose that if you know that you will soon take doxastic attitude A to *heads* as a result of *rationally* responding to new information *without loss of information*, then you should *now* take attitude A to *heads*. (This is a generalization of Bas van Fraassen's (1984) Reflection principle)" (178).[28]

However, *Value Reflection* poses a problem for fans of imprecise credences. It is generally accepted that rational imprecise agents can have credences that *dilate*, or become less precise over time.[29] But *Value Reflection* forbids imprecise agents from anticipating rational dilation. For illustration, consider the following case from White 2010:

> *Coin Game.* You haven't a clue as to whether p. But you know that I know whether p. I agree to write 'p' on one side of a fair coin, and '$\neg p$' on the other, *with whichever one is true going on the heads side* (I paint over the coin so that you can't see which sides are *heads* and *tails*). We toss the coin and observe that it happens to land on 'p'. (175)

Suppose that at the start of the coin game, you have credence $(0, 1)$ in p. If the coin lands on 'p', some members of your representor will take this outcome as confirming *heads* while others will take it as disconfirming *heads*, so after looking at the coin, you will come to have credence $(0, 1)$ in *heads*. The same thing will happen if the coin lands on '$\neg p$'. Since you know all this ahead of time, your earlier credences are a counterexample to *Value Reflection*:

$$R_0[heads|R_1[heads] = (0, 1)] = \{.5\} \neq (0, 1)$$

This counterexample has none of the usual trappings of traditional problems for Reflection. We can stipulate that you are certain that you will not forget information or update on false information, for instance, and that you are certain that you will update your representor by conditionalizing each member of it on the proposition that you learn. As a result, one might be tempted to conclude that dilation is irrational, and that the same goes for the imprecise credences that license this diachronic behavior.

27. Recall from section 2.2 that $R[p]$ is defined as $\{m(p) : m \in R\}$. An imprecise conditional credence is defined as follows: $R[p|q] =_{df} \{m(p|q) : m \in R\}$.
28. For similar remarks, see Schoenfield 2012 and Topey 2012.
29. For an early discussion of the rational permissibility of dilation, see Seidenfeld & Wasserman 1993.

How should fans of imprecise credences solve this problem? The correct diagnosis does not involve rejecting the permissibility of your imprecise priors, nor the updating rule that results in their dilation.[30] Rather, we should reject the principle of *Value Reflection* itself. At first glance, this principle may appear to be an uncontroversial extension of *Precise Reflection*. But as I shall argue, *Value Reflection* is much stronger than any principle supported by the normative facts that ground *Precise Reflection*. The appropriate extension of *Precise Reflection* is a significantly weaker principle.

The idea of tempering *Value Reflection* in response to dilation examples is discussed briefly by Schoenfield 2012 and Topey 2012. Both authors consider something like the following substitute for *Value Reflection*:

Identity Reflection $\quad R_0[p|R_1 = X] = X[p]$

Unlike *Value Reflection*, the principle of *Identity Reflection* is not violated by your *Coin Game* credences. Before you look at the coin, you have an imprecise conditional credence in *heads*, conditional on the hypothesis that you will later learn that the coin lands on 'p' and adjust your credences accordingly. The same goes for learning that the coin lands on '$\neg p$'. Conditional on either complete hypothesis about your later belief state, your current credence in *heads* is $(0, 1)$, matching the credence that you anticipate later assigning to *heads*.

Although *Identity Reflection* is consistent with the permissibility of dilation, one might worry that this version of Reflection is overly restricted in scope. Both Schoenfield and Topey raise this concern:

> We don't want the principles that tell us how to defer to be applicable only in cases where we know what the expert's entire representor is, since we rarely have such information. (Schoenfield 2012, 207)

> It isn't the case that any psychological difference renders Reflection inapplicable. If it were, Reflection would be applicable only when a person had acquired perfect knowledge of her entire future credal state. And no one ever has such knowledge. So, if Reflection is to be at all useful as a principle, some psychological differences must be irrelevant to its applicability. (Topey 2012, 485)

It is true that an imprecise agent hardly ever knows what representor she will have at a later time. But fortunately, this fact does not severely restrict the force of *Identity Reflection*. Just like *Precise Reflection*, the principle of *Identity Reflection* imposes significant rational constraints on your credences, even when you are not certain of your future credal states. *Identity Reflection* imposes constraints on your conditional credences, conditional on hypotheses about various states that you think that you might

30. Bradley & Steele 2014 investigate alternative rules for updating imprecise credences and conclude that no reasonable rule forbids rational dilation.

be in later. These constraints on your conditional credences indirectly constrain your current unconditional credences. For example, suppose that you don't know what your later representor will be, but you have .5 credence that your representor will be Q and .5 credence that it will be R, where $Q[p] = (.1, .2)$ and $R[p] = (.2, .3)$. From *Identity Reflection*, it follows that you are rationally required to believe that p is more than .15 likely and less than .25 likely, which is indeed a substantive constraint on your current credences.[31]

That being said, there is something right about the spirit of the complaints raised by Schoenfield and Topey. Fans of imprecise credences should value Reflection principles that are easy to operationalize. *Identity Reflection* constrains your credences in light of your opinions about extremely strong hypotheses. A more valuable Reflection principle would constrain your current credences in light of more targeted opinions about your future credences—for instance, constraining your current imprecise credence in p in light of your estimates of your future imprecise credence in that same proposition. Can we find a Reflection principle of this sort?

As we search for such a principle, it is again useful to pay attention to the distinction between global and pointwise constraints. So far, we have seen that global constraints can help solve several problems that cannot be solved with pointwise constraints. At this point, though, the tables have turned. In extant discussions of Reflection principles for imprecise agents, it is generally assumed that the correct analog of *Precise Reflection* for imprecise agents will be a global requirement of rationality. For example, the principle of *Value Reflection* corresponds to a global constraint on your current credences—namely, the set of representors that treat your future self as an expert about the likelihood of every proposition. In order to treat your future self as an expert about a proposition, your credences might be required to spread out to fill a certain interval, for instance, and this does not amount to the satisfaction of any pointwise constraint. Fans of imprecise credences should part with this trend in the literature, endorsing a pointwise Reflection principle. Notice that the principle of *Precise Reflection* is intimately connected the procedure of updating by conditionalization, as *Precise Reflection* requires you to defer to your future credences when you are certain you will conditionalize. Conditionalization is generally extended to imprecise agents as a rule that targets the individual elements of a representor, saying that your later representor must contain just those functions that result from conditionalizing some member of your representor on the information you learn.[32] An accompanying

31. See claim 3 of the appendix.

32. To be precise, the result of updating representor R on proposition p is $\{m(\cdot|p) : m \in R \text{ and } m(p) > 0\}$. This rule is defended using a Dutch book argument in §6.4 of Walley 1991. See Grove & Halpern 1998 and Pires 2002 for further sympathetic discussion. For a broader survey of rules for updating

Reflection principle for imprecise agents should similarly target the individual elements of a representor. Here is the rough idea behind the correct Reflection principle for imprecise agents: each individual member of your representor should defer to her own later credences, as long as she is certain that you will update rationally.

This rough idea cannot be implemented as it stands. It can be useful to talk as if your representor members are independent agents, but this metaphor has its limits. An individual member of your representor does not really have opinions about what her own later credences will be. At any given time, your total belief state is represented by a set of credence functions. There are no further facts about the cross-temporal identity of members of this set. To extend the metaphor, an individual member of your representor can only ever learn that her later credence in p will be contained in a certain set, so to speak—namely, your later imprecise credence in p.

Fortunately, this modest constraint on individual members of your representor is strong enough to yield a significant constraint on your imprecise credences. For instance, consider the fact that as a precise agent, you violate *Precise Reflection* if you have .8 conditional credence in p, given the proposition that your later credence in p is contained in $(.6, .7)$. Similarly, as an imprecise agent, you violate an important Reflection principle if some member of your representor has .8 conditional credence in p, given the proposition that your later imprecise credence in p is $(.6, .7)$. This idea is captured by the following general Reflection principle for imprecise agents:

Pointwise Reflection $R_0[p|R_1[p] = S] \subseteq S$

In other words, every probability measure m in your current representor R_0 must be such that $m(p|R_1[p] = S) \in S$. By contrast with other imprecise Reflection principles considered in this paper, *Pointwise Reflection* is a pointwise constraint, imposing a universal condition on the probability measures in your representor.

Fortunately, just like *Identity Reflection*, the principle of *Pointwise Reflection* is consistent with rational dilation. In *Coin Game*, you have .5 conditional credence in *heads*, conditional on the proposition that you will later have $(0, 1)$ credence in *heads*. This is consistent with *Pointwise Reflection*, since the former credence is contained in the latter. Meanwhile, unlike *Identity Reflection*, the principle of *Pointwise Reflection* is easy to operationalize. The principle does not constrain your credences conditional on extremely strong hypotheses about your future representor, but rather on hypotheses about your later credence in one particular proposition. Moreover, *Pointwise Reflection* is a strong rational requirement. For instance, in ordinary cases of learning where your credences shrink rather than dilate, *Pointwise Reflection* ensures that evidence

imprecise credences, see Gilboa & Schmeidler 1993.

about your later credences has a significant impact on your current credences. In the special case where S is a singleton and your representor is non-empty, *Pointwise Reflection* entails the intuitive rule of *Value Reflection* considered at the start of this section: $R_0[p|R_1[p] = S] = S$.

Pointwise Reflection is a natural generalization of *Precise Reflection*. As mentioned at the start of this section, Briggs 2009 derives a qualified version of *Precise Reflection* from the probability axioms, given modest background assumptions.[33] As long as you are certain that you will rationally update on veridical evidence, simply having a coherent credence function will guarantee that you satisfy *Precise Reflection*. We can derive a similarly qualified version of *Pointwise Reflection* from the same modest background assumptions. As long as you are certain that you will rationally update on veridical evidence, merely having a representor of coherent credence functions will guarantee that you satisfy *Pointwise Reflection*.[34] This result justifies the principle of *Pointwise Reflection* as a legitimate Reflection principle for imprecise agents, by contrast with overly ambitious principles such as *Value Reflection*.

To sum up, although it is initially tempting to extend *Precise Reflection* to an extremely strong global requirement, the spirit of this principle is best captured by a pointwise requirement. Distinguishing pointwise requirements from global requirements has thus proven useful in both directions. Articulating global independence and inertia constraints helps us solve some challenging problems for fans of imprecise credences. Articulating a pointwise Reflection constraint has helped us solve another. The notion of a global constraint is not only theoretically interesting, but also significant for the development and defense of the epistemology of imprecise credences.

33. See p. 186 of Weisberg 2007 for a similar result.
34. See claim 4 of the appendix.

Appendix

Claim 1. If $R \subseteq \{P|P(AB) = P(A)P(B)\}$, R is convex only if either $|R[A]| = 1$ or $|R[B]| = 1$.[35]

Proof: Assume R is convex. Let $P_1, P_2 \in R$. By convexity, $.5P_1 + .5P_2 \in R$. Hence:

$$\begin{aligned}
.5(P_1(A) + P_2(A)).5(P_1(B) + P_2(B)) &= .5(P_1(AB) + P_2(AB)) \\
P_1(A)P_1(B) + P_2(A)P_1(B) + P_1(A)P_2(B) + P_2(A)P_2(B) &= 2P_1(AB) + 2P_2(AB) \\
P_1(AB) + P_2(A)P_1(B) + P_1(A)P_2(B) + P_2(AB) &= 2P_1(AB) + 2P_2(AB) \\
P_2(A)P_1(B) + P_1(A)P_2(B) &= P_1(AB) + P_2(AB) \\
P_2(A)P_1(B) + P_1(A)P_2(B) &= P_1(A)P_1(B) + P_2(A)P_2(B) \\
P_2(A)[P_1(B) - P_2(B)] &= P_1(A)[P_1(B) - P_2(B)] \\
P_2(A) &= P_1(A)
\end{aligned}$$

Since P_1 and P_2 are arbitrary elements of R, it follows that $|R[A]| = 1$, as desired.

Claim 2. $C_{open} =_{df} \{R : R \text{ is not inert with respect to } \{Q : Q[H] \subseteq [.5, 1]\}\}$ is global.[36]

Proof: It suffices to produce R_1, R_2 such that $R_1 \notin C_{open}$, $R_2 \in C_{open}$, and $R_1 \subseteq R_2$. As Vallinder 2018 demonstrates, for any given interval (c_-, c_+), there exists some representor that is inert with respect to assigning that imprecise credence to H. Let R_1 be a representor that is inert with respect to assigning $(.5, 1)$ to H, and let R_2 be the union of this representor and any singleton set containing a credence function that assigns less than .5 to H. Since $R_2 \notin \{Q : Q[H] \subseteq [.5, 1]\}$, it follows that R_2 is not inert with respect to this constraint.

Claim 3. If $R_0[R_1 = Q] = .5$, $R_0[R_1 = R] = .5$, $Q[A] = (.1, .2)$, and $R[A] = (.2, .3)$, then by *Identity Reflection*, $R_0[A] \subseteq (.15, .25)$.

Proof: Let m be an arbitrary element of R_0. By the probability calculus, we have:

$$m(A) = m(R_1 = Q)m(A|R_1 = Q) + m(R_1 = R)m(A|R_1 = R)$$

35. This claim is mentioned in footnote 12 of Joyce 2010, but Joyce does not provide a proof of it.

36. I adopt the simplifying assumption that the elements of a representor are given by the expected value of a probability density function over possible chance hypotheses about the outcome of flipping a coin, where H is the proposition that the coin lands heads. As mentioned in footnote 17, the inertness relation is relative to a class of evidence propositions that is implicitly restricted by context. In this case, the evidence propositions are potential outcomes of a series of observed flips of the coin.

$m(A) = .5a + .5b$ for some $a \in (.1, .2)$, $b \in (.2, .3)$

$m(A) \in (.15, .25)$.

Claim 4. Assume that the evidence propositions that the agent might learn form a finite partition $\mathcal{B}$, and that the agent is certain that conditionalization is the right updating procedure.[37] Assume that the following qualifications hold for every $B \in \mathcal{B}$:

i. $R_0[R_0[A|B] = R_1[A|B]] = 1$

ii.[38] $R_0[B \equiv R_1[B] = 1] = 1$

iii. $R_0[R_0[A|B] = S] = 1$ if and only if $R_0[A|B] = S$

Then $R_0[A|R_1[A] = S] \subseteq S$.

Proof: The argument closely follows the derivation of Qualified Reflection on p. 69 of Briggs 2009. Let m_0 be an arbitrary element of R_0. Then we have:

$$\begin{aligned} m_0(A|R_1[A] = S) &= \frac{m_0(A \wedge R_1[A] = S)}{m_0(R_1[A] = S)} \\ &= \frac{\sum_{B \in \mathcal{B}: R_1[A|B] = S} m_0(A \wedge R_1[B] = 1)}{\sum_{B \in \mathcal{B}: R_1[A|B] = S} m_0(R_1[B] = 1)} \\ &= \frac{\sum_{B \in \mathcal{B}: R_1[A|B] = S} m_0(A|R_1[B] = 1) m_0(R_1[B] = 1)}{\sum_{B \in \mathcal{B}: R_1[A|B] = S} m_0(R_1[B] = 1)} \\ &= \frac{\sum_{B \in \mathcal{B}: R_1[A|B] = S} m_0(A|B) m_0(B)}{\sum_{B \in \mathcal{B}: R_1[A|B] = S} m_0(B)} \end{aligned}$$

By (i) and (iii), we have $R_0[A|B] = R_1[A|B]$ for every $B \in \mathcal{B}$. Hence we have:

$$\sum_{B \in \mathcal{B}: R_1[A|B] = S} m_0(A|B) = \sum_{B \in \mathcal{B}: R_0[A|B] = S} m_0(A|B)$$

By the definition of $R_0[A|B]$, we can conclude that $m_0(A|B) \in S$, and hence that $m_0(A|R_1[A] = S) \in S$. Since m_0 was chosen to be an arbitrary element of R_0, it follows that $R_0[A|R_1[A] = S] \subseteq S$, as desired.

37. The latter assumption mimics the third idealizing assumption introduced on p. 69 of Briggs 2009. According to my best understanding of Briggs, this assumption is intended to be interpreted in a way that licenses the second step of their derivation of Qualified Reflection. In order to preserve the close analogy between our arguments, I shall use the same language for the assumption that licenses the second step of my derivation of *Pointwise Reflection*.

38. This premise is inspired by premise (ii) on p. 69 of Briggs 2009, which Briggs informally glosses as the claim that the agent is "certain that she will update on veridical evidence" (70). Although Briggs formally defines premise (ii) as the claim that $Cr_0(B|Cr_1(B) = 1) = 1$, it seems to me that their derivation of Qualified Reflection requires a stronger claim, namely that $Cr_0(B \equiv Cr_1(B) = 1) = 1$.

References

Bovens, Luc. 1995. "'*P* and I Will Believe that not-*P*': Diachronic Constraints on Rational Belief." *Mind*, vol. 104 (416): 737–60.

Bradley, Seamus. 2012. "Scientific Uncertainty and Decision Making." PhD dissertation, London School of Economics and Political Science, University of London.

Bradley, Seamus & Katie Steele. 2014. "Uncertainty, Learning, and the 'Problem' of Dilation." *Erkenntnis*, vol. 79 (6): 1287–1303.

—. 2016. "Can Free Evidence Be Bad?: Value of Information for the Imprecise Probabilist." *Philosophy of Science*, vol. 83 (1): 1–28.

Briggs, Rachael. 2009. "Distorted Reflection." *Philosophical Review*, vol. 118 (1): 59–85.

Champollion, Lucas. 2020. "Distributivity, Collectivity, and Cumulativity." Forthcoming in *The Wiley Blackwell Companion to Semantics*, Lisa Matthewson, Cécile Meier, Hotze Rullmann, and Thomas Ede Zimmermann, editors. Wiley, Hoboken, NJ.

Christensen, David. 1991. "Clever Bookies and Coherent Beliefs." *Philosophical Review*, vol. 100 (2): 229–47.

Couso, Inés, Serafín Moral & Peter Walley. 2000. "A Survey of Concepts of Independence for Imprecise Probabilities." *Risk, Decision and Policy*, vol. 5 (2): 165–81.

Cozman, Fabio. 2012. "Sets of Probability Distributions, Independence, and Convexity." *Synthese*, vol. 186 (2): 577–600.

Easwaran, Kenny. 2014. "Regularity and Hyperreal Credences." *Philosophical Review*, vol. 123 (1): 1–41.

Edwards, Ward, Harold Lindman & Leonard Savage. 1963. "Bayesian Statistical Inference for Psychological Research." *Psychological Review*, vol. 70 (3): 193–242.

van Fraassen, Bas C. 1984. "Belief and the Will." *Journal of Philosophy*, vol. 81 (5): 235–56.

—. 1990. "Figures in a Probability Landscape." In *Truth or Consequences: Essays in Honor of Nuel Belnap*, J. M. Dunn & A. Gupta, editors, 345–56. Kluwer, Dordrecht.

Friedman, Jane. 2013. "Suspended Judgment." *Philosophical Studies*, vol. 162 (2): 165–81.

Gilboa, Itzhak & David Schmeidler. 1989. "Maxmin Expected Utility With Non-Unique Prior." *Journal of Mathematical Economics*, vol. 18 (2): 141–53.

—. 1993. "Updating Ambiguous Beliefs." *Journal of Economic Theory*, vol. 59 (1): 33–49.

Grove, Adam & Joseph Halpern. 1998. "Updating Sets of Probabilities." In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, Gregory Cooper & Serafin Moral, editors, 173–82. Morgan Kaufmann, San Francisco.

Hájek, Alan. 2019. "Interpretations of Probability." Available at `https://plato.stanford.edu/archives/fall2019/entries/probability-interpret/`.

Hintikka, Jaakko. 1962. *Knowledge and Belief: An Introduction to the Logic of Two Notions*. Cornell University Press, Ithaca.

Jeffrey, Richard C. 1987. "Indefinite Probability Judgment: A Reply to Levi." *Philosophy of Science*, vol. 54 (4): 586–91.

Joyce, James. 2010. "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives*, vol. 24 (1): 281–323.

Kaplan, Mark. 2010. "In Defense of Modest Probabilism." *Synthese*, vol. 176 (1): 41–55.

Kyburg, H. E. & M. Pittarelli. 1996. "Set-Based Bayesianism." *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 26 (3): 324–39.

Levi, Isaac. 1974. "On Indeterminate Probabilities." *Journal of Philosophy*, vol. 71 (13): 391–418.

—. 1980. *The Enterprise of Knowledge*. MIT Press, Cambridge.

Lewis, David K. 1980. "A Subjectivist's Guide to Objective Chance." In Lewis (1986b), 83–132. With postscript.

—. 1981. "Causal Decision Theory." In Lewis (1986b), 305–39. With postscript.

—. 1986a. *On the Plurality of Worlds*. Blackwell Publishers Inc., Malden, MA.

—. 1986b. *Philosophical Papers*, vol. 2. Oxford University Press, Oxford.

Maher, Patrick. 1992. "Diachronic Rationality." *Philosophy of Science*, vol. 59 (1): 120–41.

Moss, Sarah. 2015. "Time Slice Epistemology and Action Under Indeterminacy." In *Oxford Studies in Epistemology*, Tamar Szabó Gendler & John Hawthorne, editors, vol. 5, 172–94. Oxford University Press, Oxford.

—. 2018. *Probabilistic Knowledge*. Oxford University Press, Oxford.

Pires, Cesaltina Pacheco. 2002. "A Rule for Updating Ambiguous Beliefs." *Theory and Decision*, vol. 53 (2): 137–52.

Quine, W. V. O. 1951. "Two Dogmas of Empiricism." In *From a Logical Point of View*, 20–46. Harvard University Press, Cambridge, second edn.

Rinard, Susanna. 2013. "Against Radical Credal Imprecision." *Thought*, vol. 2 (2): 157–65.

Schoenfield, Miriam. 2012. "Chilling Out on Epistemic Rationality: A Defense of Imprecise Credences." *Philosophical Studies*, vol. 158 (2): 197–219.

Seidenfeld, Teddy & Larry Wasserman. 1993. "Dilation for Sets of Probabilities." *The Annals of Statistics*, vol. 21 (3): 1139–54.

Singer, Daniel J. 2019. "Permissible Epistemic Trade-Offs." *Australasian Journal of Philosophy*, vol. 97 (2): 281–93.

Smith, Cedric. 1961. "Consistency in Statistical Inference and Decision." *Journal of the Royal Statistical Society Series B*, vol. 23 (1): 1–37.

Stalnaker, Robert. 1984. *Inquiry*. MIT Press, Cambridge.

Talbott, William. 1991. "Two Principles of Bayesian Epistemology." *Philosophical Studies*, vol. 62 (2): 135–50.

Topey, Brett. 2012. "Coin Flips, Credences, and the Reflection Principle." *Analysis*, vol. 72 (3): 478–88.

Vallinder, Aron. 2018. "Imprecise Bayesianism and Global Belief Inertia." *British Journal for the Philosophy of Science*, vol. 69 (4): 1205–30.

Walley, Peter. 1991. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London.

Weatherson, Brian. 2008. "The Bayesian and the Dogmatist." *Proceedings of the Aristotelian Society*, vol. 107 (1 pt 2): 169–85.

—. 2015. "For Bayesians, Rational Modesty Requires Imprecision." *Ergo*, vol. 2 (20): 1–17.

Weisberg, Jonathan. 2007. "Conditionalization, Reflection, and Self-Knowledge." *Philosophical Studies*, vol. 135 (2): 179–97.

White, Roger. 2010. "Evidential Symmetry and Mushy Credence." In *Oxford Studies in Epistemology*, Tamar Szabó Gendler & John Hawthorne, editors, vol. 3, 161–86. Oxford University Press, Oxford.

Williams, J. Robert G. 2014. "Decision Making Under Indeterminacy." *Philosophers' Imprint*, vol. 14 (4): 1–34.

Williams, P. M. 1976. "Indeterminate Probabilities." In *Formal Methods in the Methodology of Empirical Sciences*, M. Przelecki, K. Szaniawski & R. Wojcicki, editors, 229–46. Reidel, Dordrecht.