

Uncertainty, learning, and the “problem” of dilation

Seamus Bradley and Katie Steele

August 11, 2013

Abstract

Imprecise probabilism – which holds that rational belief/credence is permissibly represented by a set of probability functions – apparently suffers from a problem known as *dilation*. We explore whether this problem can be avoided or mitigated by one of the following strategies: a) modifying the rule by which the credal state is updated, b) restricting the domain of reasonable credal states to those that preclude dilation.

1 Introduction

Imprecise probabilism – the view that your belief or credal state is best represented by a *set of probability functions* – has received a lot of attention recently. One important prima facie problem with imprecise probabilism is *dilation*. This is a puzzling phenomenon whereby in certain conditions, when you update on a piece of evidence, your degrees of belief get *less precise*.

In Section 2, we initially introduce imprecise probabilism and give an example of the phenomenon of dilation. Then we give a more formal treatment of our assumed rule for updating imprecise probabilities. This in turn allows a clearer understanding of what the phenomenon of dilation is and when it happens. With these preliminaries in place, Section 3 investigates alternative belief-updating rules that handle imprecision, and considers desiderata for an appropriate update rule. The upshot of this discussion is that no reasonable update rule avoids dilation. Section 4 changes tack: we consider whether dilation can be avoided by restricting the domain of rational prior credal states. We discuss why these sorts of restrictions are not convincing, and go on to conclude that credal states that result in dilation may in fact be reasonable.

2 Imprecise probabilism and dilation

We need some preliminaries before we can properly engage with a discussion of dilation and imprecise probabilism. Not least, we need definitions of probabilism, imprecise probabilism and dilation.

2.1 Probabilism, precise and imprecise

Orthodox probabilism has your credal state represented by a probability function, which is a function Pr that maps *events*, X , to the unit interval, such that:

- $\text{Pr}(X) \geq 0$
- If the X_i are mutually exclusive and exhaustive, $\sum \text{Pr}(X_i) = 1$

We say that $\text{Pr}(X)$ is your degree of belief in the event X .

Imprecise probabilism has your credal state represented by a *set* of such functions called \mathcal{P} . We can then think of $\mathcal{P}(X)$ as the set of values assigned to the proposition X by $\text{Pr} \in \mathcal{P}$. It can be taken to be a representation of your unsharp degree of belief in X . There are a number of arguments for representing belief in this way, not least that it hardly seems a *requirement* of rationality that belief be precise (and preferences complete); surely imprecise belief (and corresponding incomplete preferences) are at least rationally *permissible*.

One important positive motivation for imprecise probabilism is to represent the difference between the *weight* of evidence and the *balance* of evidence (Joyce, 2005). For example, if you have a coin that you have not tossed at all, and you have no reason to think heads or tails more likely, but nothing to rule out arbitrary bias in either direction either, then the *balance* of evidence suggests that the coin is as likely as not to land heads. That is, no evidence pushes in one direction or the other. Now imagine tossing the coin a hundred times and observing about fifty heads. The evidence is still balanced evenly between heads and tails, but now there is more *weight* behind the claim that it is as likely as not to land heads. The salient precise probabilist response to both examples is to assign degree of belief 0.5 to the event “the coin will land heads”. But there seems to be an important difference between the weight of evidence in these cases. The imprecise probabilist can represent this difference by assigning $[0, 1]$ to the heads event in the first case and $\{0.5\}$ in the second.¹

¹See also Sturgeon (2008) and Kaplan (2010).

2.2 Dilation, informally

Imagine you believe that there are a total of 10 black and 10 white marbles distributed somehow among the urns X and Y . Each urn contains 10 marbles. An urn will be selected at random by flipping a fair coin, and a marble drawn from it. Using X and Y to refer to the propositions “The marble is drawn from urn X ” and “The marble is drawn from urn Y ” respectively; and using B and W to stand for the propositions “The marble drawn is black” and “The marble drawn is white” respectively, the following is a plausible characteristic of your belief *representor*,² \mathcal{P} : $\Pr(B|X) = 1 - \Pr(B|Y)$ for all $\Pr \in \mathcal{P}$. As such, before learning, your credences regarding the colour of the marble drawn are as follows: $\mathcal{P}(B) = \{0.5\} = \mathcal{P}(W)$. You believe that the number of white marbles and the number of black marbles are equal and that over the two urns their probabilities average out.

It is surely plausible that your conditional credences are, however, imprecise: you have no information about how the marbles are distributed between the urns, and so it is plausible that you do not rule out the possibility that urn X contains only white marbles and also the possibility that X contains no white marbles, and everything in between. Thus your representor \mathcal{P} plausibly includes probability functions that represent each of those possibilities. That is, $\mathcal{P}(W|X) = \mathcal{P}(B|X) = \mathcal{P}(W|Y) = \mathcal{P}(B|Y) = \{0, \frac{1}{10}, \dots, \frac{9}{10}, 1\}$, or, if *convexity* were mandated (as per Levi (1974, 1986)), your representor would plausibly have $\mathcal{P}(W|X) = [0, 1]$, and likewise for the other conditional attitudes. Whether convexity *is* mandated when in fact the probability of drawing white given X could not possibly be, say, $\frac{6}{21}$, is a tricky question and one we shall ignore. Nothing in our discussion hinges on the sets of probabilities being convex. For reasons of notational convenience we will stick to using $[0, 1]$ to represent the aforesaid imprecise conditional beliefs. So learning which urn is drawn from *dilates* your probability for white from $\{0.5\}$ to $[0, 1]$ and likewise for black. That is, your beliefs in B and W get *less precise*, once you learn either X or Y . Note that this analysis of the problem makes an assumption about what the right belief after learning ought to be. It assumes that after learning X , say, $\{\Pr(\cdot|X) : \Pr \in \mathcal{P}\}$ is your new representor.

Note that the standard examples for dilation in the literature involve odd correlations, such as between the outcome of a coin toss and the truth of some event for which evidence is lacking (White (2010); Pedersen and Wheeler (ms); Joyce (2011); Walley (1991)). These cases tend to generate confusion about proper application of the Principal Principle (see Lewis (1986)); it is not immediately obvious what is “admissible” and what is “inadmissible”

²The term “representor” is due to van Fraassen (1990).

evidence for, say, the coin-toss outcome. Such confusion gets in the way of intuitions regarding dilation and makes analysis of the phenomenon more difficult. We prefer the example above because it makes the case for dilation more straightforward.

Dilation is taken to be a serious problem for imprecise probabilism. Dilation seems to lead to violation of prominent epistemic principles like *reflection*; it also looks to have unpalatable decision theoretic consequences; furthermore it just seems odd that evidence can make beliefs less precise.³ The aim of this paper is to find out how the phenomenon of dilation arises, and whether it can or should be avoided, at least from the epistemic point of view (we save the decision-theoretic point of view for another occasion).

2.3 Updating imprecise probabilities

In this section we borrow some formalism from Grove and Halpern (1998) in order to clarify the standard rule for updating imprecise beliefs, and the conditions that generate dilation. We have a set of possible worlds S and an algebra defined on them \mathcal{E} . This is the space of possibilities we are considering: $\mathcal{M} = \langle S, \mathcal{E} \rangle$. There is a set of probability functions defined on \mathcal{M} which we denote $\Pi^{\mathcal{M}}$. We think of the process of updating as a function Upd that takes as input a set of probabilities \mathcal{P} and an event, and outputs a set of probabilities. That is: $\text{Upd}: 2^{\Pi^{\mathcal{M}}} \times \mathcal{E} \rightarrow 2^{\Pi^{\mathcal{M}}}$.

Recall that $\mathcal{P}(X)$ is the set of values assigned to X by elements of \mathcal{P} . The following definitions will be useful:

- $\underline{\mathcal{P}}(X) = \inf\{\text{Pr}(X) : \text{Pr} \in \mathcal{P}\}$
- $\overline{\mathcal{P}}(X) = \sup\{\text{Pr}(X) : \text{Pr} \in \mathcal{P}\}$
- $\underline{\mathbb{P}}(X) = \{\text{Pr} \in \mathcal{P} : \text{Pr}(X) = \underline{\mathcal{P}}(X)\}$
- $\overline{\mathbb{P}}(X) = \{\text{Pr} \in \mathcal{P} : \text{Pr}(X) = \overline{\mathcal{P}}(X)\}$

$\underline{\mathcal{P}}(X)$ is the infimum of $\mathcal{P}(X)$ and $\overline{\mathcal{P}}(X)$ the supremum. $\underline{\mathcal{P}}$ and $\overline{\mathcal{P}}$, thought of as real valued functions, behave like lower and upper probabilities respectively (Walley, 1991; Cozman, nd; Halpern, 2003). Colloquially speaking, $\underline{\mathcal{P}}(X)$ is the lowest value assigned to X , and $\underline{\mathbb{P}}(X)$ is the set of probability functions that assign that value to X . Likewise for $\overline{\mathcal{P}}$ and $\overline{\mathbb{P}}$.

Generalised conditioning is the standard rule for updating sets of probabilities. This is denoted Upd_{cond} . Generalised conditioning is characterised

³This consequence also seems to be in conflict with the weight/balance motivation for imprecise probabilism mentioned earlier.

as follows:

$$\text{Upd}_{\text{cond}}(\mathcal{P}, E) = \{\text{Pr}(\cdot|E), \text{Pr} \in \mathcal{P}, \text{Pr}(E) \neq 0\} \quad (1)$$

This is the rule we assumed above, and it is subject to dilation. It is not the only possible method for updating, but it does seem like the most natural generalisation of Bayesian conditioning to the imprecise setting.

2.4 Dilation, formally

We say an update rule Upd is subject to dilation for representor \mathcal{P} when there exist $X, E \in \mathcal{E}$ such that:

$$\underline{\text{Upd}}(\mathcal{P}, E)(X) < \underline{\mathcal{P}}(X) \leq \overline{\mathcal{P}}(X) < \overline{\text{Upd}}(\mathcal{P}, E)(X)$$

Given an update rule Upd and a representor \mathcal{P} , if the above holds, we say E dilates X .⁴ $\text{Upd}(\mathcal{P}, E)$, which, recall, denotes the general updating function, is a set of probabilities, so it makes sense to think about its extrema. In short, the above expressions are meaningful.

When is dilation possible? The typical characterisation of dilation found in e.g. Seidenfeld and Wasserman (1993) and Pedersen and Wheeler (ms) is in terms of divergence from stochastic independence. This characterisation assumes that updating works by Upd_{cond} . So the following conditions are not universal. But for a broad class of rules that work in roughly the same way as Upd_{cond} – classical rules (see later) – these conditions are telling.

Define $S(X, E) = \text{Pr}(X \cap E) / \text{Pr}(X) \text{Pr}(E)$. This can be understood as a measure of the correlation of X and E . Now define:

- $\mathbb{S}^+(X, E) = \{\text{Pr} \in \mathcal{P}, S(X, E) > 1\}$
- $\mathbb{S}^-(X, E) = \{\text{Pr} \in \mathcal{P}, S(X, E) < 1\}$

$\mathbb{S}^+(X, E)$ is the set of probability functions in the representor that have X and E positively correlated. $\mathbb{S}^-(X, E)$ is the set of probability functions that have X and E negatively correlated.

A necessary condition for dilation is:

$$\underline{\mathbb{P}}(X|E) \subseteq \mathbb{S}^-(X, E) \text{ and } \overline{\mathbb{P}}(X|E) \subseteq \mathbb{S}^+(X, E) \quad (2)$$

If E dilates X , then the above holds. That is, it is necessary for dilation that the updated lower probability is such that the events in question were

⁴Throughout the paper it will normally be obvious what rule and representor we are discussing, and so we do not always make this explicit when describing a case of dilation.

negatively correlated; and the updated upper probability is such that the events were positively correlated. Consider the opposite scenario: if the lower probability for X after learning E were not associated with probability functions that had X and E negatively correlated, then there would not have been a *decrease* in lower probability for X . Likewise for the upper end.

A sufficient condition for dilation is:

$$\underline{\mathbb{P}}(X) \cap \mathbb{S}^-(X, E) \neq \emptyset \text{ and } \overline{\mathbb{P}}(X) \cap \mathbb{S}^+(X, E) \neq \emptyset \quad (3)$$

That is, if the above holds, then E dilates X . It is sufficient for dilation that there is some probability function which has a minimal probability for X prior to learning and which has X and E negatively correlated; and that there is some probability function which has a maximal probability for X prior to learning and has X and E positively correlated. One can see this by thinking of \mathcal{P} as a *credal committee*.⁵ Each probability function in \mathcal{P} stands for a committee member with particular opinions. If there are committee members who think X more unlikely than does anyone else and who also think E makes X even less likely; and there are committee members who think X more likely than does anyone else and who also think E makes X even more likely, then learning E will have the effect of moving the opinions of those members further apart: dilation.

We later refer to the above conditions – one necessary, the other sufficient – when we try to find ways to block dilation. The main reason for presenting these conditions, however, is that they allow a clearer understanding of the phenomenon of dilation.

Often a distinction is drawn between dilation and *strict dilation*, which is when every element E_i of some partition dilates belief in the one proposition X . Strict dilation is the focus of the literature to date, or more precisely, special symmetrical cases of strict dilation (like our urns example) are the focus of the literature. These special cases of dilation are considered troubling because they apparently raise a conflict for general principles like reflection that many find intuitively compelling (as we will see in Sections 4.2 and 4.3). But we hold that non-symmetrical cases of strict dilation and standard dilation are puzzling too. Imprecise probabilities are supposed to allow us to represent the distinction between weight of evidence and balance of evidence; a simple reading of this claim would have it that more evidence (greater weight of evidence) means more precision. Dilation seems to be a case where more evidence leads to less precision. Therefore, dilation – strict or otherwise – problematizes this popular line of argument.

⁵Joyce (2011) attributes this term to Adam Elga.

3 Classical update rules

This section investigates whether some plausible alternative to generalised conditioning may preclude dilation. We restrict our attention to an important class of update rules known as *classical rules*⁶ that have this form:

$$\text{Upd}(\mathcal{P}, E) = \{\text{Pr}(\cdot|E), \text{Pr} \in \mathcal{P}', \mathcal{P}' \subseteq \mathcal{P}\} \quad (4)$$

for some suitably chosen \mathcal{P}' . That is, the updated belief is some subset of the set of conditional probabilities. We do not defend classical updating rules in this paper; we simply note that classical rules have the best claim to being natural generalisations of standard Bayesian conditioning, which serves as our benchmark for updating rules in the precise context. Standard Bayesian conditioning moves from a prior probability to the relevant conditional probability. It is natural to think that a set of probabilities is updated by moving to some set of conditional probabilities. Classical rules capture this intuition.⁷ Note that the classical rules are the ones for which the characterisation of dilation mentioned above is relevant.

The largest such classical rule is generalised conditioning. In the case of generalised conditioning we take \mathcal{P}' to be $\{\text{Pr} \in \mathcal{P} : \text{Pr}(E) > 0\}$. Other choices for \mathcal{P}' that also rest on stipulations regarding $\text{Pr}(E)$ are:⁸

1. $\{\text{Pr} \in \mathcal{P} : \text{Pr}(E) = \overline{\text{Pr}}(E)\}$ (in other words, $\overline{\mathbb{P}}$)
2. $\{\text{Pr} \in \mathcal{P} : \text{Pr}(E) > \tau\}$ for some fixed $\tau < 1$
3. $\{\text{Pr} \in \mathcal{P} : \text{Pr}(E) = 1\}$
4. \emptyset

Rules of this form typically have one of two problems: either they suffer from dilation (the top two); or they are “empty” – i.e. $\text{Upd}(\mathcal{P}, E) = \emptyset$ – when we would want them not to be (the bottom three).

In any case, the “maximum likelihood” rule (number 1 on the list above) and its cousins seem to miss the point if the aim is to avoid dilation. There is no reasonable restriction for \mathcal{P}' based on $\text{Pr}(E)$ such that updating on E never causes dilation. Consider the marbles-in-urns case again. Dilation

⁶The terminology is from Gilboa and Schmeidler (1993).

⁷We won’t discuss more radical departures from the standard Bayesian model, such as Kyburg’s Evidential Probabilities model (Kyburg and Teng, 2001). Kyburg’s model doesn’t really have a concept of updating: there is simply the rationally permissible belief function given a certain body of evidence.

⁸The first of these is the rule that Gilboa and Schmeidler (1993) endorse. The third is mentioned, but not endorsed by Grove and Halpern (1998).

occurs whatever the (non-extreme) probability of picking urn X . It is most striking for $\Pr(X) = 0.5$, but any other value for $\Pr(X)$ besides zero or one also has the same problem.⁹ So the prior probability of the evidence does not seem like the right kind of restriction on updating if we are interested in avoiding dilation. We should, instead, be looking for some restriction that speaks to our characterisation of dilation in Section 2.

So what we want is some method of selecting a subset of $\text{Upd}_{\text{cond}}(\mathcal{P}, E)$ such that this updated representor is never subject to dilation. Let's try to formulate such a classical update rule now. We shall use the shorthand \mathcal{P}_E to mean $\text{Upd}(\mathcal{P}, E)$. What an update rule needs to do to block dilation is the following: Whenever $\underline{\mathcal{P}}_E(X) < \underline{\mathcal{P}}(X)$ and $\overline{\mathcal{P}}_E(X) > \overline{\mathcal{P}}(X)$ at least one of the following two sets of probability functions must be excised from the posterior representor:

$$\begin{aligned}\underline{Q} &= \{\Pr \in \mathcal{P}, \Pr(X|E) \in [\underline{\mathcal{P}}_E(X), \underline{\mathcal{P}}(X)]\} \\ \overline{Q} &= \{\Pr \in \mathcal{P}, \Pr(X|E) \in [\overline{\mathcal{P}}(X), \overline{\mathcal{P}}_E(X)]\}\end{aligned}$$

That is, whenever dilation would otherwise occur, the updated representor should contain no elements of \underline{Q} or it should contain no elements of \overline{Q} .

This is the property that a classical rule must satisfy in order to block dilation. It is a very direct approach. What we are basically saying is “whenever the set of probability functions after updating would have resulted in dilation, you should remove those updated probability functions that yield the dilation.” This is an extremely *ad hoc* property; it is a very artificial way of avoiding dilation.

The dilation-blocker property also has worrying implications. Consider an example like the one above with marbles in urns, but now let's say you know that all the black marbles are in one urn, and all the white marbles in the other. You just don't know which is which. That is, your representor contains just the two endpoints of the representor considered above: the two functions with $\Pr(B|X) = 1$ and $\Pr(B|X) = 0$. So now learning that the urn drawn from is X causes your representor to dilate to the two values for B : $\text{Upd}(\mathcal{P}, X)(B) = \{0, 1\}$. To block dilation, one or the other of these probability functions must be removed from the representor. In either case, you become certain of what colour marble will be drawn from the urn. This is unintuitive, and at odds with the aim of imprecise probabilities, which is to represent radical uncertainty.

One might argue that this is an artefact of having a non-convex representor. If you took a convex cover of the above representor, this would

⁹The updated belief for B is still $[0, 1]$, whatever value $\mathcal{P}(X)$ takes. For values of $\mathcal{P}(X)$ not equal to $\{0.5\}$, the prior belief for B is imprecise.

contain $\Pr(B|X) = 0.5$, and indeed $\text{Upd}(\mathcal{P}, X)(B) = \{0.5\}$ seems like the only plausible classical update that would avoid dilation in this example. We do not want to enter the debate about whether convexity is mandated.¹⁰ We will, however, point out that the non-convex representor gets *something* right about the epistemic state after learning X : it is now *determined* what colour the marble will be, even if you don't know which way it is determined.

Let us nonetheless restrict our attention for the moment to convex representors, and reflect further on these marbles-in-urns examples. Take the case where the probability for “the marble drawn is black” (B) dilates from $\{0.5\}$ to $[0, 1]$ according to generalised conditioning. How should an alternative update rule that avoids dilation deal with this case? As mentioned above, it seems that the only reasonable possibility is for the updated probability to remain $\{0.5\}$. No other possible update would respect the symmetry of the situation. But now it seems like you have become convinced that X and B are probabilistically independent. This is not something you were sure of before updating.

More generally, avoiding dilation in the urns examples requires that $\mathcal{P}_X(B)$ ends up being a subset either of $[0, 0.5]$ or of $[0.5, 1]$. Therefore, you end up becoming certain of whether B and X are positively or negatively correlated. Again, this does not seem to be in keeping with the motivation for imprecise probabilism. In short, removing probability functions from the updated representor does not seem like a good approach to imprecise updating. Removing functions with $\Pr(E) = 0$ from \mathcal{P}_E is acceptable, because those functions are defective – they assigned zero probability to an event that actually happened. Any other sort of removal does not seem warranted.

Things are in fact worse than this. However we block dilation, one of \overline{Q} or \underline{Q} will be removed. This means that there are $\Pr \in \mathcal{P}$ such that $\Pr(\cdot|E)$ is not in \mathcal{P}_E for the relevant E . This immediately leads to worries that such an update rule will violate *commutativity*: if some other event, F , say, were learned first, before E , then the aforementioned \Pr 's that would be excised from the representor if E were learned first may well not yield dilation if E is learned second, and thus not be excised from the final updated representor. Indeed, any rule that removes some probability functions from the updated set suggests worries about commutativity. Generalised conditioning respects commutativity because it only removes those \Pr with $\Pr(E) = 0$, namely those \Pr shown to be incompatible with the evidence; these \Pr will not ultimately be live options anyway.

The upshot of this discussion is that, *if the domain of prior representors*

¹⁰Recall from Section 2.2 our ambivalence about whether convexity is mandated. But see Kyburg and Pittarelli (1992) for discussion of this issue.

is unrestricted,¹¹ then no classical update rule avoids dilation while also adequately representing the post-update belief state and satisfying the desirable constraint of commutativity. Given that generalised conditioning does answer to the latter two concerns – it tracks the full extent of uncertainty in the course of learning and is commutative – we will appeal to this belief update rule in the remainder of the paper. To reiterate our reasoning: given that dilation cannot be avoided with an alternative classical update rule, we will stick with the most popular such rule for updating imprecise beliefs, namely generalised conditioning. Note that generalised conditioning is axiomatised by Grove and Halpern (1998), and can also be defended by way of a Dutch book argument (Walley, 1991, Section 6.4).¹² So we think it warranted to assume this particular classical rule in the remainder of the paper.

4 Return of the prior

It might be argued that the problem of dilation lies not in the update rule but rather in your prior representor. Maybe by appropriately restricting what sort of priors the update rule should reasonably have to deal with, we can block dilation. That is, if we could argue that “dilation-vulnerable” priors were irrational, *in any situation*, then there would be no problem. In other words, this section investigates whether there should be a restriction on what prior representors are permissible, such that all representors that are vulnerable to dilation are ruled out as impermissible.

In what follows, we examine two proposals to this effect that appear in some form in the literature. They should be understood as norms of belief that are more fundamental than the details of the formal representation of belief. The proposals have much in common, and similar arguments apply to them. The first proposal (4.2) explores a norm relating to irrelevance that precludes certain structures i.e., certain relations between beliefs, in your prior representor. The second proposal (4.3) exploits a version of the reflection principle to the same effect – certain relations between present and anticipated future beliefs are ruled out.

In these two sections we consider just whether the proposed norms are compelling, and whether they serve to rule out the salient dilation-vulnerable priors in our original marbles-in-urns example. There are serious doubts as

¹¹We return to restrictions on the priors in the next section.

¹²Recall that classical rules have credence match up, in some sense, with your conditional beliefs. Walley shows that, if these prior conditional attitudes are what govern your conditional betting behaviour, then generalised conditioning is the right sort of update rule to avoid the possibility of a Dutch book being made against you.

to whether the proposals are plausible at all, and, even if plausible, whether they do in fact succeed in ruling out the prior representor at issue in our urns example. In any case, there is a stronger reason for claiming that these proposals do not serve to block dilation: at best, they only preclude a certain kind of “symmetrical” dilation. The proposals permit many other cases of even strict dilation. We discuss this in Section 4.4.

Before we turn to the main content of this section, we would like to clarify how our project differs from another project, which seeks to rule out cases of dilation as cases where the formal model misdescribes certain intuitive features of the set up. This we do in Section 4.1.

4.1 Irrelevance and model-building

The odd thing about dilation in our urns example is that you are learning something – X , say – that does not discriminate between a black or white marble being drawn from the urn, and so seems *irrelevant* to the beliefs it dilates. That is, learning which urn is drawn from (X or Y) does not seem relevant to the colour of the ball drawn from the urn (B or W). So why should learning X , say, have such a disastrous effect on your belief in B ? Surely something is amiss here.

In the standard precise setting, irrelevance or independence of X and B is typically modelled by: $\Pr(B|X) = \Pr(B)$, $\Pr(X|B) = \Pr(X)$ or by $\Pr(BX) = \Pr(B)\Pr(X)$. In this setting these three expressions are equivalent. Call the first epistemic irrelevance of X to B , the second epistemic irrelevance of B to X , and the third stochastic independence of B and X .

These distinctions are important in the imprecise case since analogues of these properties are not equivalent. For sets of probabilities Pedersen and Wheeler (ms) give the analogue of epistemic irrelevance of X to B as:¹³

1. $\underline{\mathcal{P}}(B|X) = \underline{\mathcal{P}}(B) = \underline{\mathcal{P}}(B|\neg X)$
2. $\overline{\mathcal{P}}(B|X) = \overline{\mathcal{P}}(B) = \overline{\mathcal{P}}(B|\neg X)$

It should be clear that if your representor satisfied this condition, X would not dilate B . A stronger formal property is that of *epistemic independence*, which holds for X and B if and only if X is epistemically irrelevant to B and B is epistemically irrelevant to X . The imprecise analogue of stochastic

¹³Note that $\mathcal{P}(B|X)$ is a prior conditional belief, not an updated belief. It is the set of values assigned to B conditional on X by the set of *prior* probabilities \mathcal{P} . That is, $\mathcal{P}(B|X) = \{\Pr(B|X) : \Pr \in \mathcal{P}\}$. This is extensionally the same as the updated belief in B after updating on X (assuming generalised conditioning), but we vary the formalism to highlight the conceptual distinction between updated belief and conditional belief.

independence is stronger again: X and B are stochastically independent if and only if, for every $\text{Pr} \in \mathcal{P}$, we have $\text{Pr}(BX) = \text{Pr}(B)\text{Pr}(X)$; in other words, when stochastic independence (in the precise sense) holds for every probability function in your representor.

The question now becomes: is it reasonable to demand that even the weakest of these notions of irrelevance holds (that is, epistemic irrelevance of X to B) for the formal model of our marbles-in-urns example? This property does not hold in our original model of the situation as described in Section 2.2: we have $\underline{\mathcal{P}}(B|X) = 0$ but $\underline{\mathcal{P}}(B) = 0.5$. But perhaps this reflects the fact that our initial formal rendering of the problem does not capture the problem at hand, precisely because it does not have X irrelevant to B . There are two ways we could modify the prior representor in our example to satisfy this condition. One way would be to make $\underline{\mathcal{P}}(B) = 0$ and $\overline{\mathcal{P}}(B) = 1$; the other, to make $\underline{\mathcal{P}}(B|X) = 0.5$, and make commensurate changes to the other conditional probabilities. The first of these is not plausible, since we are taking $\mathcal{P}(B) = \{0.5\}$ to be non-negotiable: it is fixed by the problem set-up (via the Principal Principle). So it is implausible to change $\mathcal{P}(B)$ and thus implausible to change $\underline{\mathcal{P}}(B)$.

What about building in epistemic irrelevance by making $\underline{\mathcal{P}}(B|X) = 0.5$? Can we find a model of the marbles-in-urns example that is plausible, but is such that $\underline{\mathcal{P}}(B|X) = 0.5$? Recall that in the original model $\underline{\mathcal{P}}(B|X) = 0$. So relative to the original model, our new model would involve removing some Pr from our initial description of \mathcal{P} . $\underline{\mathcal{P}}(B|X) = 0.5$ entails that for all $\text{Pr} \in \mathcal{P}$, $\text{Pr}(B|X) \geq 0.5$; to also have $\overline{\mathcal{P}}(B|X) = 0.5$ requires that $\text{Pr}(B|X) \leq 0.5$ for all $\text{Pr} \in \mathcal{P}$. Likewise we require that $\overline{\mathcal{P}}(B|Y) = 0.5 = \underline{\mathcal{P}}(B|Y)$. So this is a case where your credal committee is in agreement that it is equally likely that a marble be drawn from urn X as from urn Y . But this is at odds with the imprecision in the description of the problem: the whole point of the example is that you *don't* know whether a black marble is more likely from urn X or urn Y . The reasoning is much the same as in Section 3: removing probability functions from the updated representor doesn't do justice to the imprecision in the problem set-up. So it is not plausible to impose the epistemic irrelevance of X to B in your prior representor in this way either.

Our original model of the set-up does seem to be the most faithful to the informal description we gave of it. At the very least, the model described in Section 2.2 is surely legitimate. Our intuition that X is somehow irrelevant to B is not borne out by the model. We can reconcile ourselves to this idea as follows: learning X *seems* irrelevant to B because it is evidence of unknown value. It is not that X is not relevant to B , it's just that you don't know how relevant X is to B . There is a real but unknown correlation between B and X . This is a case of what Pedersen and Wheeler (ms) call *proper dilation* as

opposed to improper dilation where the dilation can be mitigated by paying attention to epistemic irrelevancies among the events.¹⁴ So imposing the epistemic irrelevance of X to B in our example is not justified, at least not without more of a story. We turn to such a story now.

4.2 An irrelevance norm for belief?

There may yet be something to the intuition that learning X should not radically alter your belief in B . The way to make this compelling is to articulate a norm of belief that serves as a clear constraint on rational belief models. In other words, we seek a plausible norm of belief that effectively restricts the space of *rationally permissible* belief representors, such that the belief representor described in Section 2.2 would be deemed irrational. Some representors are such that they are vulnerable to dilation: that is, there exist propositions such that the conditions for dilation described in Section 2.4 are satisfied. Other representors are not vulnerable to dilation. If there were some plausible constraint on rational belief representors that would make all members of the former category impermissible, then dilation could be ruled out. This is so since no permissible rational belief would be vulnerable to dilation.¹⁵

Drawing inspiration from (Joyce, 2011, p. 302), we propose the following, which we call the irrelevance norm.¹⁶

If it is the case that your belief in some proposition B supposing X is the same as your belief in B supposing $\neg X$,¹⁷ then it must also be the case that your unconditional belief in B is the same as the conditional beliefs.

¹⁴Pedersen and Wheeler (ms) explore a number of cases where the initial model is somehow misdescribed, and once the appropriate redescription is done, certain troubling instances of dilation disappear. These are the cases of improper dilation; where the initial model does not take into account certain kinds of irrelevance that should be built into the model. Our example is not such a case.

¹⁵Earlier we implicitly assumed that the domain of the update rule was universal: Upd was a function on all of $2^{\Pi^{\mathcal{M}}}$. Now we are considering restricting Upd to some subset of $2^{\Pi^{\mathcal{M}}}$. The hope is that there is some plausible restriction that precludes Upd being vulnerable to dilation.

¹⁶Note that Joyce does not offer a norm of belief, but rather a definition of epistemic (or “evidential” irrelevance) that is based on the pattern of likelihoods across some evidence partition. Joyce’s definition of irrelevance is different from those discussed by Pedersen and Wheeler (ms). Our norm amounts to roughly the following restriction: “If X is irrelevant to B (in Joyce’s sense) then X should be epistemically irrelevant to B (in the sense of Pedersen and Wheeler)”.

¹⁷Of course, beliefs given X and beliefs given $\neg X$ will differ with respect to, say, the truth of X . But what is required here is that they amount to the same beliefs *about* B .

When rendered as a constraint on formal representations of belief, this condition can be read as:

If it is the case that $\mathcal{P}(B|X) = \mathcal{P}(B|\neg X)$ then it must also be the case that $\mathcal{P}(B) = \mathcal{P}(B|X) = \mathcal{P}(B|\neg X)$

The idea is that if learning X is going to change your belief in B , it shouldn't change it in the same way that learning $\neg X$ would.¹⁸ The salient imprecise beliefs in our urns example (Section 2.2) do not conform to the above norm. We have $\mathcal{P}(B) = \{0.5\}$, by the Principal Principle. This is non-negotiable. The conditional beliefs for B , given X and given $\neg X$ (i.e. Y) are as follows: $\mathcal{P}(B|X) = \mathcal{P}(B|\neg X) = [0, 1]$. These conditional beliefs are identical, which means, according to the irrelevance norm, that belief in B should be the same as belief in B given X . That is, it should be the case that the aforesaid conditional beliefs are equivalent to the unconditional belief in B , which is $\{0.5\}$. But this is not the case. Overall then, this particular arrangement of beliefs is illegitimate according to the irrelevance norm. Moreover, due to symmetry considerations and $\mathcal{P}(B) = \{0.5\}$ being non-negotiable, the only plausible and legitimate beliefs are ones that have $\mathcal{P}(B|X) = \mathcal{P}(B|\neg X) = \{0.5\}$. That is, the rational agent is effectively forced to have precise conditional attitudes in our urns example, and dilation is thereby precluded.

Of course, the question remains: is the irrelevance norm a reasonable constraint on rational belief? The imprecise enthusiast would surely object to restrictions on the domain of permissible prior representors. In the case of the urns example, the irrelevance norm, coupled with the non-negotiable instance of the Principal Principle requires you to regard two events such as B and X as not only epistemically independent, but also stochastically independent, when the natural way to think of these events is as having unknown correlation. The points from the previous sections carry some weight here too. The worry is that this precise model of the urns example doesn't do justice to the definite but unknown correlation between B and X .

One may of course conclude that the irrelevance norm is *not* a reasonable norm for rational belief. Interestingly, Joyce does not seem to take this route, but rather emphasises that complementary beliefs are not to be confused with identical beliefs. Considered as sets of values, $\mathcal{P}(B|X)$ and $\mathcal{P}(B|\neg X)$ are identical, but what Joyce is arguing is that *the credal states they represent are distinct*. When we look more closely, we find that these conditional beliefs in fact differ, because for each probability function Pr in your representor, where $\text{Pr}(B|X) = p$, we see that $\text{Pr}(B|\neg X) = 1 - p$. So Joyce might

¹⁸This is not a mathematical truth, of course, but rather a substantial restriction on rational belief functions. This is why we describe it as a *norm* of belief.

accept the informal version of the irrelevance norm, but deny that the formal rendering we gave is adequate. Note that this defense effectively highlights a deficiency in the standard formalism for discussing imprecise belief. Joyce would say that the irrelevance norm is not violated in our urns example due to complementarities between the urns: the conditional attitudes upon learning X or $\neg X$ are not identical, but complementary, so there is no need for the prior belief in B to be the same.

We can accept the nuanced version of the irrelevance norm (à la Joyce), or we can reject the norm outright; either way, the prior representor allowing dilation remains permissible for our urns example. If, on the other hand, we accept the formal rendering of the irrelevance norm, then dilation is ruled out, at least for the case considered up until now. We return to whether such a norm rules out all cases of dilation in Section 4.4.

4.3 A reflection norm for belief?

Reflection is the idea that, if you know that later, your beliefs will be thus and so, and you consider them rationally justified, then you ought to believe now what you expect yourself to believe in the future.¹⁹ Some consider this principle to be an appropriate restriction on rational belief, and moreover argue that the dilating imprecise beliefs in cases like our urns example fall afoul of the principle (see White (2010), cf. Topey (2012)).

Note that White takes the violation of reflection (a principle that he clearly holds dear) in examples akin to our urns example to be a refutation of imprecise probabilism. In effect, White assumes that imprecise probabilism involves a commitment to a universal domain of priors; i.e. imprecise probabilism is committed to any constellation of imprecise beliefs that fits the basic structure (sets of probability functions) being permissible. Since various examples make vivid that, given anticipated future learning, some constellations of imprecise beliefs violate reflection as stated above, White concludes that imprecise probabilism must be rejected.

Here we investigate a proposal that is similar to White's, but which does not involve a commitment to universal domain for imprecise probabilism. As per our discussion above, the question is whether there is a plausible restriction on rational belief that may in effect rule out (all) dilation-vulnerable

¹⁹This is a limited version of the *reflection principle*. It applies just to cases where you *know* what your beliefs will be in the future, presumably because, whatever the evidence, you will have the same beliefs. We do not deny that more general versions of reflection are interesting and may constrain rational belief, but in the context of our discussion, the limited version of the principle is more pertinent and is already controversial.

priors. The *reflection principle* is a plausible candidate for such a principle, and perhaps more palatable than the above-discussed irrelevance norm.

Consider our urns example, this time in the light of the reflection principle: assuming that your beliefs are as we described them above, you know that your belief in B will be $[0, 1]$ after rational updating, since both possible events you could learn, X and $\neg X$ (Y), lead to the $[0, 1]$ update via generalised conditioning. So according to reflection, you ought to believe that now. That is, your prior belief in B ought to be $[0, 1]$. But your prior belief in B is $\{0.5\}$ (by the Principal Principle). So this constellation of prior beliefs is illegitimate. Moreover, symmetry considerations coupled with reflection suggest that the only plausible beliefs one can have here are as follows: $\mathcal{P}(B) = \mathcal{P}(B|X) = \mathcal{P}(B|\neg X) = \{0.5\}$.

The reader may well anticipate how the counter-arguments go. Staunch defenders of imprecise probabilism who are committed to the universal domain claim may simply reject the reflection principle. They may claim that while reflection seems intuitive, it is not in fact a fundamental principle of rational belief, and cases like our urns example simply demonstrate why reflection does not always hold. Otherwise, one could take the Joyce line and argue that reflection is a legitimate constraint on rational belief, but despite first appearances, it does not in fact apply to the urns case. The point of contention – as in the last section – is whether the future beliefs are identical. If the beliefs in B on learning X and on learning $\neg X$ are identical, but different from the prior belief in B , then reflection rules out this collection of beliefs as irrational. But the updated beliefs are *not* identical in this example, even though they are both represented in summary form as $[0, 1]$. The updated beliefs are rather complementary. Or so the argument goes.

As with the case of the irrelevance norm, we think the jury is still out on this issue of whether, in the context of principles like reflection, imprecise conditional beliefs of $[0, 1]$ should be regarded as the *same belief*, regardless of any complementarities.²⁰ In any case, neither the irrelevance norm nor reflection will serve to block all dilation-vulnerable priors, even if they block dilation-vulnerable priors in highly symmetrical cases like the urns example. We turn to this issue of generality now.

4.4 The search for a general “dilation blocker”

The following example serves to demonstrate why neither an irrelevance-based norm nor a reflection-based norm is sufficiently general to rule out all

²⁰Toppey (2012) offers some interesting reasons why conditional beliefs of $[0, 1]$ *should* be treated as identical, and thus why the reflection principle *does* apply in cases like our urns example.

dilation-vulnerable priors. Indeed such norms cannot even rule out all cases of strict dilation. In short, there are cases of strict dilation (and of course standard dilation) that are not symmetrical. By symmetrical we mean that $\mathcal{P}(X|E) = \mathcal{P}(X|\neg E)$. Consider the following case: There are four possible “states” arising from the product of two independent binary events: $\{E, \neg E\} \times \{H, T\}$. Assume that $\{H, T\}$ is the outcome of a fair coin toss, so your beliefs are $\mathcal{P}(H) = \mathcal{P}(T) = \{0.5\}$. Your beliefs about E are indeterminate, such that $\mathcal{P}(E) = [0.1, 0.6]$ and $\mathcal{P}(\neg E) = [0.4, 0.9]$. Define the event $F = \{EH, \neg ET\}$. Some calculation reveals that your prior belief in F is $\{0.5\}$. Now let us consider what would happen to your belief in F , were you to learn H or T . We get $\mathcal{P}(F|H) = [0.1, 0.6]$, and $\mathcal{P}(F|T) = [0.4, 0.9]$.²¹ That is, we have a case of strict dilation for F , given the evidence partition $\{H, T\}$. It is clear, however, that neither the irrelevance norm nor the reflection principle can preclude this case of strict dilation, as your conditional beliefs after learning are obviously not identical. But F dilates H . So even if we had been successful in blocking some cases of dilation by ruling out some collections of priors with irrelevance or reflection principles, these principles wouldn’t speak to this case of dilation.

One may well maintain that the irrelevance norm and the reflection principle are compelling restrictions on prior belief functions, and moreover, that they should be interpreted in a way that actually does restrict the class of prior representors. This would be to say that these principles are more fundamental, as it were, than universal domain, the principle that any set of probability functions is a legitimate, rationally permissible prior belief. In our concluding remarks below, we are more sympathetic to the universal domain idea. But the point stressed in this section is that no interpretation of the irrelevance norm or the reflection norm discussed here can serve to preclude all cases of dilation-vulnerable priors.

A direct way to block all cases of dilation is to propose a principle that requires prior belief functions do not anywhere satisfy the necessary conditions for dilation. This would be akin to our “dilation-blocker” update rule, except that on this proposal, the restrictions apply to your prior belief function; updating would be unproblematic because you would not be permitted the dilation-vulnerable priors to begin with. Any such principle is hardly acceptable, however. As per the “dilation-blocker” update rule, the principle is extremely ad hoc, as it has no independent motivation (unlike irrelevance or reflection).

²¹We owe this example of non-symmetric dilation to Teddy Seidenfeld (in correspondence).

5 Concluding remarks

We have seen that dilation is a problem that plagues many kinds of classical rules for updating imprecise probabilities. The classical rules that avoid dilation have other problems, like non-commutativity, or empty update sets. Indeed, if we want to avoid these sorts of problems then generalised conditioning is the only plausible update rule. But this rule evidently does not preclude dilation when the domain of prior representors is unrestricted.

The further question is whether an unrestricted domain of prior representors is reasonable. Perhaps prior representors that lead to dilation (via generalised conditioning) are simply irrational. We explored a couple of proposals to this effect – the irrelevance principle and the reflection principle – but found them to be controversial, and in any case, not sufficiently general to “block” all cases of dilation.

Let us close with a final suggestion that is admittedly deflationary about the *epistemic* puzzle of dilation. The suggestion is simple, and is taken for granted in some circles, but we trust the journey to this position presented in this paper is nonetheless worthwhile.²² We suggest that dilation, with respect to the transition from, say, $\mathcal{P}(B)$ to $\mathcal{P}_E(B)$, is not as epistemically puzzling as it first appears. It is not puzzling because the conditions for dilation are already present in your prior representor, before learning. For instance, it must be the case that the necessary condition for dilation is fulfilled in your prior conditional credences. The dilation is not a new and unexpected event that occurs upon learning. It is just a portion of your prior representor that becomes realised upon learning.²³ That is, in the same way that a suitably reflective agent with precise beliefs can predict her own updated belief in X on learning E (i.e. $\Pr(X|E)$) for not-yet observed E , a suitably reflective imprecise agent will see that her credal state will dilate. The dilation-vulnerable parts of your prior conditional beliefs simply indicate cases where there is evidence of unknown value to be learned.²⁴

It is also wrong to say that your prior representor \mathcal{P} is “precise” and your posterior representor is “imprecise”. As it happens, \mathcal{P} assigns a sharp value to the events B and W , but other events have imprecise prior belief: consider

²²The position we describe here seems to be the consensus view in statistics. Our aim is to present that view to philosophy.

²³Walley (1991) suggests something in this vein as an attempt to reconcile his readers to dilation.

²⁴Note that this deflationary suggestion about dilation does nothing to rehabilitate the “weight of evidence” motivation for imprecise probabilism that we mentioned earlier. This motivation for imprecise probabilism may well be misguided, but note that there are other motivations that imprecise probabilists can appeal to. See, for instance, Joyce (2011).

$\mathcal{P}(BX)$, for example. What we draw attention to is that your epistemic situation with respect to B , say, is expressed by your *entire belief representor*, or at least some relevant portion of your representor that includes conditional beliefs about B , and not just $\mathcal{P}(B)$ on its own. The complementarities Joyce draws attention to are a reason to be sceptical that $\mathcal{P}(B)$ is an adequate representation of your belief in B . Moreover, there are other reasons to think that this representation is deficient: consider your degree of belief that a coin of unknown bias will land heads. Plausibly, $\mathcal{P}(H) = [0, 1]$. Now, consider your degree of belief that the coin will land heads ten times in a row: $\mathcal{P}(10H) = [0, 1]$. Do you have the same attitude to $10H$ as you do to H ? Surely not! Whatever the real chance of heads, you know that H will be at least as likely as $10H$, and almost certainly strictly more likely. But again, the sets of probability values amounting to $\mathcal{P}(H)$ and $\mathcal{P}(10H)$ fail to reflect this fact.

Once we attend to your full suite of prior beliefs concerning B , there is no reason to be puzzled about dilation. The change of belief in question was already written into your prior beliefs. So from the epistemic point of view, dilation in unconditional belief due to learning is not such a peculiar phenomenon. In fact, the balance of considerations suggest that the imprecise probabilist is better off, from the purely epistemic point of view, to reconcile herself with dilation, rather than to opt for some side constraints on imprecise priors or update rules that serve to preclude dilation. Of course, the story does not end there: dilation can lead to somewhat undesirable consequences for sequential decision-making. The decision-theoretic investigation of dilation, however, must be left for another paper.²⁵

References

- Bradley, S. and Steele, K. (ms.a). Can free evidence be bad? value of information for the imprecise probabilist.
- Bradley, S. and Steele, K. (ms.b). Subjective probabilities need not be sharp.
- Cozman, F. (2012). Sets of probability distributions, independence and convexity. *Synthese*, 186:577–600.
- Cozman, F. (n.d.). A brief introduction to the theory of sets of probability measures. <http://www.poli.usp.br/p/fabio.cozman/Research/CredalSetsTutorial/quasi-bayesian.html>.

²⁵See Weatherson (ms), Joyce (2011) and Bradley and Steele (msa,m)

- Gilboa, I. and Schmeidler, D. (1993). Updating ambiguous beliefs. *Journal of Economic Theory*, 59:33–49.
- Grove, A. and Halpern, J. Y. (1998). Updating sets of probabilities. In *Proceedings of the fourteenth conference on uncertainty in AI*, pages 173–182.
- Halpern, J. Y. (2003). *Reasoning about uncertainty*. MIT press.
- Joyce, J. M. (2005). How probabilities reflect evidence. *Philosophical Perspectives*, 19:153–178.
- Joyce, J. M. (2011). A defense of imprecise credence. *Oxford Studies in Epistemology*, 4. forthcoming.
- Kaplan, M. (2010). In defense of modest probabilism. *Synthese*, 176:41–55.
- Kyburg, H. E. and Pittarelli, M. (1992). Set-based Bayesianism. Technical Report UR CSD;TR407, University of Rochester, Computer Science Department. <http://hdl.handle.net/1802/765>.
- Kyburg, H. E. and Teng, C. M. (2001). *Uncertain Inference*. Cambridge University Press.
- Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy*, 71:391–418.
- Levi, I. (1986). *Hard choices: decision making under unresolved conflict*. Cambridge University Press.
- Lewis, D. (1986). A subjectivist’s guide to objective chance (and postscript). In *Philosophical Papers II*, pages 83–132. Oxford University Press.
- Pedersen, A. P. and Wheeler, G. (ms.). Demystifying dilation.
- Seidenfeld, T. and Wasserman, L. (1993). Dilation for sets of probabilities. *Annals of Statistics*, 21:1139–1154.
- Sturgeon, S. (2008). Reason and the grain of belief. *Noûs*, 42:139–165.
- Topey, B. (2012). Coin flips, credences and the reflection principle. *Analysis*, 72:478–488.
- van Fraassen, B. (1990). Figures in a probability landscape. In Dunn, M. and Segerberg, K., editors, *Truth or Consequences*, pages 345–356. Kluwer.

Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*, volume 42 of *Monographs on Statistics and Applied Probability*. Chapman and Hall.

Weatherson, B. (ms.). Decision making with imprecise probabilities. Available <http://brian.weatherson.org/vdt.pdf>.

White, R. (2010). Evidential symmetry and mushy credence. *Oxford Studies in Epistemology*.