

CREDAL IMPRECISION AND THE VALUE OF EVIDENCE

Nilanjan Das
University College London

Abstract

This paper is about a tension between two theses. The first is *Value of Evidence*: roughly, the thesis that it is always rational for an agent to gather and use cost-free evidence for making decisions. The second is *Rationality of Imprecision*: the thesis that an agent can be rationally required to adopt doxastic states that are *imprecise*, i.e., not representable by a single credence function. While others have noticed this tension, I offer a new diagnosis of it. I show that it arises when an agent with an imprecise doxastic state engages in an *unreflective inquiry*, an inquiry where they revise their beliefs using an updating rule that doesn't satisfy a weak reflection principle. In such an unreflective inquiry, certain synchronic norms of instrumental rationality can make it instrumentally irrational for an agent to gather and use cost-free evidence. I then go on to propose a diachronic norm of instrumental rationality that preserves *Value of Evidence* in unreflective inquiries. This, I suggest, may help us reconcile this thesis with *Rationality of Imprecision*.

Consider this argument. The success of our actions depends on the way the world is. We can only find out how the world is by gathering more evidence and revising our beliefs in light of it. So, if a piece of evidence is available for gathering and use, it always makes sense to gather that evidence and use it for making decisions.

This argument—though plausible at first glance—is subject to two challenges. First, if the expected costs of gathering and using the available evidence outweigh the expected benefits of doing so, then it doesn't make sense to gather and use that evidence. Second, if an agent is uncertain about their own rationality or is misinformed about the requirements of rationality, then it may not be rational for them to gather and use cost-free evidence. So, we may qualify the conclusion of our argument as follows.

Value of Evidence. Suppose a piece of evidence is available to an agent for gathering and use at a negligible cost, and the agent is rationally and correctly certain that they will respond to that evidence in an epistemically and instrumentally rational manner. Then, it is instrumentally rational for that agent to gather that evidence and use it for making decisions.

Traditionally, philosophers have defended this principle by relying on the tools of decision theory.¹ But it turns out that such decision-theoretic arguments depend on a

¹See Pierce (1967), Ramsey (1990), Blackwell (1951) and Good (1967).

number of non-trivial background assumptions. When we relax these assumptions, the arguments for *Value of Evidence* fail.

Here, I will focus on one such assumption: namely, that the relevant agent's doxastic state should be *precise*, i.e., their doxastic attitude towards any proposition should be representable by a single real number that reflects their *credence* or *degree of belief* in that proposition. Many people reject this assumption. They favour:

Rationality of Imprecision. In some cases, an agent may be required by epistemic rationality to adopt an imprecise doxastic state.²

Several writers have noticed that, given some background assumptions about rationality, a tension emerges between *Value of Evidence* and *Rationality of Imprecision*.³ If these assumptions are right, then an agent with an imprecise doxastic state may find themselves in a case of *dilation*, a case where their doxastic attitude towards a proposition rationally goes from being precise to being imprecise (i.e., from being representable by a precise real number to being representable only by a non-singleton set of real numbers) in response to some evidence. In these cases, certain decision rules can make it irrational for such an agent to gather and use cost-free evidence. So, *Value of Evidence* will be false.

In this paper, I will do three things. First, I will offer a diagnosis of this conflict between *Value of Evidence* and *Rationality of Imprecision*. I will show that this conflict arises when an agent with an imprecise doxastic state revises their beliefs according to a generalized version of Bayesian conditionalization and, as a result, fails to satisfy a weak reflection principle. In such cases, they engage in what I call an *unreflective inquiry* (§§1-4). In such an unreflective inquiry, a range of synchronic norms of instrumental rationality—norms according to which what an agent is permitted to do at a time doesn't directly depend on what they do (or are disposed to do) at other times—can lead to failures of *Value of Evidence* (§§5-9).

Second, I will propose a diachronic norm of instrumental rationality—called *Practical Stability*—that helps us preserve *Value of Evidence* in these cases (§10). According to *Practical Stability*, an agent should act in a diachronically stable manner in the course of any inquiry: roughly speaking, they should pick a credence function from the set that represents their doxastic state at the beginning of their inquiry and choose acts that are optimal by lights of that credence function, or rationally revised versions of that credence function, throughout their inquiry.

Third, I will go on to explain how *Practical Stability* can be defended. On the one hand, it can independently supported by means of a *Dutch book argument*, which shows that an agent who violates this norm will be predictably vulnerable to a sure loss of

²For a defence of this view in philosophy, see Levi (1974, 1980), Kaplan (1983), Joyce (2005, 2010), and Sturgeon (2008).

³See, for example, Good (1974), Grünwald and Halpern (2004), Kadane, Schervish and Seidenfeld (2008), Pedersen and Wheeler (2015), and Bradley and Steele (2016).

utility (§11). On the other hand, it can be defended against a number of salient objections (§§12-14). This, in turn, might give us hope that the tension between *Value of Evidence* and *Rationality of Imprecision* can be resolved.

1 The Rationality of Imprecision

In this section, I will motivate *Rationality of Imprecision* by appealing to the following case of *non-specific* or *incomplete* evidence.⁴

Mystery Urn 1. I am rationally certain that an opaque urn contains 10 balls that are either blue or green, and that a ball will be randomly picked from the urn. But I have no clue about what proportion of the balls are blue.

Call the proposition that the selected ball will be blue B . In this case, my evidence is non-specific or incomplete: I have no evidence about the distribution of balls in the urn. What should my doxastic attitude towards B be?

According to one answer, it should be precise; it should be representable by a single real number that reflects my credence or degree of belief in B . This answer is supported by two principles: the Principle of Indifference and Lewis' (1980) Principal Principle. In this case, there are eleven hypotheses about possible distributions of balls in the urn:

H_0 : There are ten green balls and no blue balls in the urn.

H_1 : There are nine green balls and one blue ball in the urn.

...

H_{10} : There are no green balls and ten blue balls in the urn.

Note two facts. First, the Principle of Indifference says that, for any two propositions X and Y , if an agent has no reason to treat X or Y as more likely than the other, then they should assign the same credence to each of them.⁵ Since I have no reason to treat any H_i to be more likely than any other H_j (where i and j lie between 0 and 10), I should assign an equal credence of $\frac{1}{11}$ to each H_i on the list. Second, the Principal Principle says that, given the evidence that the objective chance of a proposition X is r , an agent's credence in that proposition should be r (in the absence of inadmissible information). In this case, if any H_i is true, the objective chance of B is $\frac{i}{10}$. For, if there are i blue balls in the urn, the objective chance of a blue ball being picked is $\frac{i}{10}$. So, according to the Principal Principle, my conditional credence in B given any H_i should be $\frac{i}{10}$. Therefore, by the law of total probability, my unconditional credence in B should be $\frac{0 + 0.1 + \dots + 0.9 + 1}{11} = 0.5$.

⁴Versions of this argument have been considered by Walley (1991, p. 2), Joyce (2005, 2010) and Elga (2010).

⁵While the principle can be traced back to Laplace, it has been defended Jaynes (1957), White (2009), Hawthorne, Landes, Wallmann, and Williamson (2017), and Eva (2019).

The Principle of Indifference is unpopular: relative to distinct partitions of a possibility space, it can require an agent to assign different credences to the same proposition.⁶ Suppose we reject the Principle of Indifference on these grounds. Should I then assign some other precise credence to B ? The answer seems to be “No.” I have no reason to treat any of H_i ’s as more likely than any other. Thus, I can’t assign any other precise credence to B .

Yet, if we think that there is a unique doxastic attitude that I am required to adopt towards B ,⁷ then the only remaining option for us is to say that my doxastic attitude towards B should be imprecise: a range of real numbers lying between 0 and 1 (corresponding to all the possible assignments of credences to the H_i ’s). This will also make the overall doxastic state imprecise. It won’t be representable by a single credence function, but rather by a set of credence functions Γ , such that the set of credences assigned to B by the credence functions in Γ —call it $\Gamma(B)$ —is the interval $[0,1]$. This supports *Rationality of Imprecision*.

Following van Fraassen (1990), I will call the set of credence functions that represent an agent’s precise or imprecise doxastic state their *representor*. Defenders of *Rationality of Imprecision* suggest that we think of an agent’s representor as a *credal committee*.⁸ In the case of an agent with a precise doxastic state, the credal committee contains only one member. But, in the case of an agent with an imprecise doxastic state, the committee contains several members, each of whom has a different opinion on some matter. While the opinions of an individual committee member don’t reflect the beliefs of the agent themselves, the committee as a whole can represent the doxastic state of the agent.

Defenders of *Rationality of Imprecision* often assume that an agent’s rational representor satisfies two constraints. The first is *probabilistic coherence*: the representor of an epistemically rational agent only contains probability functions. The second is *convexity*: if p_1 and p_2 are any two probability functions in an epistemically rational agent’s representor, then any probability function p_3 that is a mixture of p_1 and p_2 is also a member of the representor.⁹ I will assume the first, but not the second.

In the next three sections (§§2-4), I will show how *Rationality of Imprecision* clears room for the possibility of unreflective inquiries.

⁶This problem—sometimes called *problem of multiple partitions*—has been discussed by Keynes (1921), van Fraassen (1989) and White (2009).

⁷We may reject this assumption if we accept *permissivism*, the view that it can be rationally permissible for two different agents, or the same agent at different times, to adopt different doxastic attitudes towards the same proposition in response to the same body of evidence. For arguments against this view, see White (2005), Horowitz (2014, 2019), Greco and Hedden (2016), and Schultheis (2018). For some arguments in favour of permissivism, see Douven (2009), Titelbaum (2010), Schoenfield (2014, 2019), and Titelbaum and Kopec (ms).

⁸Joyce (2010) ascribes this idea to Adam Elga.

⁹In other words, if Γ is the relevant representor, then, for any two probability functions p_1 and p_2 in Γ , $\alpha.p_1 + (1 - \alpha).p_2$ must be also be in Γ for any α between 0 and 1 (inclusive). Levi (1974, 1980) defends this view.

2 Unreflective Inquiries

In this section, I will lay out the concept of a *unreflective inquiry*.

Let an *inquiry* be any evidence-gathering act or event that satisfies two conditions. First, prior to that act or event, the relevant agent is (or can be) rationally certain that they will engage in that act or undergo that event. Second, as a result of that act or event, the agent couldn't lose any evidence, but might gain some evidence. I will represent any such inquiry with a triple $\langle W, E, \Gamma \rangle$. Here, W is a finite set of worlds that are compatible with the agent's *total prior evidence*, i.e., their total evidence before their inquiry.¹⁰ E is a function that maps each world w in W to a proposition (or, a set of worlds in W) that reflects the agent's *total posterior evidence* in w , i.e., their total evidence in w after the inquiry. Finally, Γ is the agent's *rational prior representor*—a non-empty set of regular probability functions defined on subsets of W —which reflects the agent's rational doxastic state before the inquiry.¹¹ So, if an agent's rational doxastic state is precise before their inquiry, then Γ will contain a single probability function. But, for an agent whose rational prior doxastic state is imprecise, Γ will contain more than one probability function. Throughout this essay, I will assume that the agent's rational prior representor is the same at every world w in W . This means that, if the agent's rational prior representor is Γ , then their prior evidence before their inquiry also entails that their rational prior representor is Γ .

In the course of an inquiry, an agent may revise their beliefs in response to new evidence according to an *updating rule*, a rule that tells the agent what doxastic state they should adopt in response to the evidence they get. We can represent an updating rule as a function U that maps each world w in W to a posterior representor, i.e., a set of probability functions defined on subsets of W , which represents a recommended doxastic state the agent could have after their inquiry. I'll assume that updating rules don't discriminate between worlds where the agent's evidence remains the same. So, for any two worlds w and w^* in W and any updating rule U , if $E(w) = E(w^*)$, then $U(w) = U(w^*)$.¹²

For the bulk of this essay, we shall be concerned with a weak reflection principle that constrains updating rules. Here's a version of that principle that applies to agents with precise doxastic states: an agent with a precise doxastic state shouldn't update according to a rule U in the context of an inquiry if, by their own lights, U is guaranteed to lower their credence in some proposition in the course of that inquiry. For example,

¹⁰I am assuming that W is finite in order to avoid discussing certain failures of *Value of Evidence* that arise in infinitary cases due to violations of conglomerability; see, for discussion of this problem, Kadane, Schervish and Seidenfeld (2008) and Pedersen and Wheeler (2015). The problem we discuss here doesn't have the same solution as this other problem.

¹¹A regular probability function in this context is a probability function p such that, for any $w \in W$, $p(\{w\}) > 0$. In the rest of this essay, I will abuse this notation and write " $p(w)$ " instead of " $p(\{w\})$."

¹²For some discussion of the same constraint on updating rules, see Greaves and Wallace (2006) and Schoenfield (2017).

imagine I rationally assign a credence of 0.5 in the proposition that I'm not popular.¹³ I am about to consult my friend about my popularity. I am rationally sure that they will tell me that I'm popular. But I'm also extremely gullible, and I am aware of this. So, I am also rationally certain that I will take my friend's testimony at face value and become certain that I am popular. So, my credence that I'm not popular will drop to 0. Here, my updating rule will violate the weak reflection principle that I've sketched above.

We can formulate the principle a bit more carefully.

Weak Reflection. For any inquiry $\langle W, E, \Gamma \rangle$, suppose that:

- Γ is a singleton set of probability functions, and
- U is an updating rule such that, for any w in W , $U(w)$ is a singleton set of probability functions.

Then, the relevant agent is permitted by epistemic rationality to update according to U only if there is no proposition X such that, for any p in Γ and any world w in W , there is some p^* in $U(w)$ for which $p^*(X) < p(X)$.

According to *Weak Reflection*, if an updating rule U is epistemically rational for an agent (with a precise doxastic state) to comply with, then there is no proposition X such that future credences in X recommended by U are uniformly lower than the agent's prior credence in X . It is motivated by the same idea that motivates other reflection principles: namely, that an epistemically rational agent should revise their beliefs according to an updating rule only if, by their own lights, their prior doxastic state coheres with the possible future doxastic states recommended by that rule.¹⁴

It is easy to generalize *Weak Reflection* to the case of agents with imprecise doxastic states. According to the generalized version of the principle, an agent should use an updating rule U in the course of an inquiry only if there isn't any proposition X such that, for any credence function p in an agent's rational prior representor, we can find a credence function p^* in any posterior representor recommended by U such that $p(X) >$

¹³Examples of this sort are discussed by Salow (2018).

¹⁴*Weak Reflection* is weaker than van Fraassen's (1984) reflection principle. According to the latter, an agent's prior credence in any proposition X —conditional on the proposition that their future credence in X recommended by their updating rule is r —should also be r . This implies that the agent's prior credence in any proposition X should be a weighted average of the possible future credences in X recommended by the relevant updating rule (where the weights are their prior credences that those will be the credences that the rule will actually recommend). So, the possible future credences in X recommended by the rule cannot be uniformly lower than the agent's prior credence in X by their own lights. Thus, this principle entails a version of our weak reflection principle. However, our weak reflection principle doesn't entail van Fraassen's principle. Take a case where an agent assigns a prior credence of 0.5 to a proposition X , but also assign non-zero prior credence to possibilities where the future credence in X (recommended by their updating rule) is either 0.4 or 0.5. Here, the relevant updating rule doesn't violate our *Weak Reflection*, since the future credences in X recommended by that rule aren't uniformly lower than their prior credence in X . But the agent's prior credence X cannot be expressed as a weighted average of 0.4 and 0.5 if the weights in question are positive. So, van Fraassen's reflection principle fails.

$p^*(X)$. To see what this means, imagine again that I have a sharp credence of 0.5 that I am unpopular. I am also rationally certain that I will soon consult two friends about my popularity: one will say that I'm unpopular, while the other will deny this. I have no clue about how reliable they are. So, a bit strangely, after speaking to them, I will assign an interval of credences $[0, 1]$ to the proposition that I'm unpopular. This is a case of dilation: if this were to happen, my doxastic attitude towards the proposition that I am unpopular would go from being precise to being imprecise. The generalized version of *Weak Reflection* says that I shouldn't update in this way. In this case, the only credence function p in my (rational) prior representor assigns a credence of 0.5 to the proposition that I am unpopular, but there are plenty of credence functions p^* in my posterior representor, which assign a credence lower than 0.5 to the same proposition. Thus, my updating rule in this case violates the generalized version of *Weak Reflection*.

We can state the generalized version of *Weak Reflection* as follows.

Generalized Weak Reflection. For any inquiry $\langle W, E, \Gamma \rangle$, an agent is permitted by epistemic rationality to update according to an updating rule U only if there is no proposition X such that, for any p in Γ and any world w in W , there is some p^* in $U(w)$ such that $p^*(X) < p(X)$.

This is exactly like *Weak Reflection*, except that we have now dropped the restriction that the agent's prior and posterior representors should be singleton sets of probability functions. Let any updating rule that conforms to this constraint be *reflective*. More specifically, for any inquiry $\langle W, E, \Gamma \rangle$, an updating rule U is reflective just in case there is no proposition X such that, for any p in Γ and any world w in W , there is some p^* in $U(w)$ such that $p^*(X) < p(X)$. Let any updating rule that isn't reflective be *unreflective*. And let an inquiry be *unreflective* just in case, in the course of that inquiry, the agent revises their doxastic states using an unreflective updating rule.¹⁵

The idea of an unreflective inquiry will prove useful for us. In the next two sections (§§3-4), I will show that, if *Rationality of Imprecision* is right, then—given a plausible rule of rational belief revision—an agent with an imprecise doxastic state will sometimes rationally engage in an unreflective inquiry.

¹⁵White (2009) also notes that, in cases of dilation, a version of van Fraassen's reflection principle will fail. According to this version of the principle, if an agent rationally certain that they will adopt a certain doxastic attitude towards a proposition X after their inquiry by rationally responding to their evidence without any loss of information, then they should adopt that doxastic attitude towards X prior to the inquiry. This would imply that I can't be rationally certain that I will rationally assign an interval of credences $[0,1]$ to the proposition that I am unpopular, while currently rationally assigning a sharp credence of 0.5 to this proposition. White's reflection principle entails the *Generalized Weak Reflection Principle* (given the assumption that the relevant agent is rationally certain of what epistemic rationality permits them to do). Moss (2020) argues such failures of reflection aren't a problem.

3 Reflection and Conditionalization

Bayesian conditionalization yields a rule of rational belief revision for agents with precise doxastic states. It says that, if an agent's credence function at a time t is p and the strongest evidence they gain between t and a later time t^* (without losing any evidence) is a proposition E_i , then their posterior credence in any proposition H at t^* should just be the conditional credence $p(H|E_i)$ (provided that $p(E_i) > 0$). Defenders of *Rationality of Imprecision* claim that an agent with an imprecise doxastic state should update their beliefs according to the following generalized version of Bayesian Conditionalization.

Generalized Conditionalization. Suppose Γ is an agent's rational representor at t , and the strongest evidence that the agent gains between t and a later time t^* (without losing any evidence) is a proposition E_i . Then, their doxastic attitude towards any proposition H at t^* should be represented by $\Gamma(H|E_i) = \{p(H|E_i) : p \in \Gamma\}$, provided that $p(E_i) > 0$ for every probability function p in Γ .¹⁶

Generalized Conditionalization makes sense in light of the "credal committee" analogy. If an agent's prior representor is a committee whose members have different opinions on some matter, then the agent's posterior representor should simply be the same committee, except that its members should by now have rationally revised their earlier opinions in light of the new evidence. If Bayesian conditionalization is the right rule of rational belief revision for each committee member, then *Generalized Conditionalization* is true.

I'll call any updating rule that conforms to *Generalized Conditionalization* a *conditionalizing rule*. For any inquiry $\langle W, E, \Gamma \rangle$, an updating rule U is a conditionalizing rule just in case, for any w , $U(w) = \{p(\cdot|E(w)) : p \in \Gamma\}$. The important question for us is this: Can conditionalizing rules be unreflective?

A few writers have shown that, at least for agents with precise doxastic states, a conditionalizing rule U can be unreflective when a condition called *partitionality* fails.¹⁷

¹⁶Some version of this principle is assumed by Kadane, Schervish and Seidenfeld (2008), Pedersen and Wheeler (2015) and Bradley and Steele (2016). One standard worry—discussed by Levi (1980, ch. 13), Walley (1991, p. 93), Joyce (2010), Rinard (2013), Vallinder (2018), and Moss (2020)—about *Generalized Conditionalization* is that it gives rise to cases of belief inertia, where an agent who initially assigns an interval of credences $[0,1]$ to a proposition cannot then rationally move to a sharper doxastic attitude. This makes inductive learning with such vacuous priors impossible. There are a few standard responses to this problem. Joyce (2010) and Rinard (2013) say that we shouldn't assign such vacuous priors to any proposition. For example, Joyce claims that an agent's representor is subject to some constraints of rationality that eliminate "pigheaded" and "extremist" elements from that representor. Vallinder (2018) argues that Joyce's proposal doesn't succeed. Moss (2020) solves the problem—satisfactorily I think—by proposing an updating rule that requires us to accommodate not only propositional evidence, but also evidence that is irreducibly probabilistic in form. So, in response to information about objective chances, some extreme members of the agent's representor may indeed drop out, and make way for inductive learning. But this doesn't create trouble for *Generalized Conditionalization*, since it only applies to cases where the strongest evidence that agent gains is a proposition.

¹⁷See Williamson (2000, ch. 10), Salow (2018), and Das (2020).

An inquiry $\langle W, E, \Gamma \rangle$ is partitional just in case the agent’s posterior total evidence has two features. First, it is *factive*: it only entails truths. So, for any world w in W , w is in $E(w)$. Second, it is *perfectly introspective*: if the agent’s posterior evidence is E_i , it entails that their posterior evidence is exactly E_i . So, for any two worlds w and w^* in W , if w^* is in $E(w)$, $E(w) = E(w^*)$. Together, these two features imply that, if E_1, E_2, \dots, E_k are the strongest pieces of posterior evidence that the agent could get in the course of their inquiry, then these evidence propositions form a partition over W such that the agent’s posterior evidence in any E_i -world is E_i . But if the agent’s posterior evidence isn’t factive or doesn’t have perfect introspective access to itself, a conditionalizing rule can indeed end up being unreflective.

As Das (2020) shows, since failures of partitionality can give rise to unreflective inquiries, they can lead to failures of *Value of Evidence*. But this phenomenon has nothing much to do with *Rationality of Imprecision*. So, we shall set aside non-partitional inquiries. Thus, our question becomes: Can conditionalizing rules be unreflective in the case of partitional inquiries? For agents with precise doxastic states, the answer is “No.”¹⁸ But, for agents with imprecise doxastic states, it is “Yes.” In the next section, I will describe an example where a conditionalizing rule is unreflective even though the inquiry is partitional.

4 An Unreflective Inquiry

The example, once again, is a case of dilation.¹⁹

Mystery Urn 2. There are two opaque urns, H and T . I am rationally certain that H and T together contain twenty balls, ten of which are blue while the rest are green. But I have no clue about how many balls in each urn are blue. I am also rationally certain that a fair coin has been flipped: if the coin landed heads, a ball will be randomly drawn from H , and, if it landed tails, a ball will be randomly drawn from T . I am about to learn what the outcome of the coin flip was.²⁰

Let B be the proposition that the selected ball is blue. Let *Heads* be the proposition that the coin has landed heads, and let *Tails* be the proposition that the coin has landed tails. This inquiry is partitional. The strongest pieces of evidence that I could get in the

¹⁸This is easy to check. Suppose for *reductio* that there is a partitional inquiry $\langle W, E, \Gamma \rangle$ such that Γ contains a single probability function and U is a conditionalizing rule that is unreflective. Suppose $\Gamma = \{p\}$. Since U is a conditionalizing rule, for any w , $U(w) = \{p(\cdot|E(w))\}$. Since the inquiry is partitional, for any X (that is a subset of W), $p(X) = \sum_{w \in W} p(w)p(X|E(w))$. But, since U is unreflective, for every w , $p(X) > p(X|E(w))$. These two claims are incompatible. So, if U is a conditionalizing rule, then either the inquiry isn’t partitional or U isn’t unreflective.

¹⁹Some writers such as White (2009) argue that dilation is problematic. Hart and Titelbaum (2015) argue that the case that White relies on is independently questionable, while Pedersen and Wheeler (2014) distinguish good cases of dilation from bad ones.

²⁰The example is adapted from Bradley and Steele (2016).

course of my inquiry are *Heads* and *Tails*. These form a partition over the possibility space, such that, in any *Heads*-world, I learn *Heads* and, in any *Tails*-world, I learn *Tails*. We can show that, in this case, before I learn about the outcome of the coin flip, my rational doxastic attitude towards *B* should be precise: every member of my rational prior representor will assign a credence of 0.5 to *B*. But, if I update according to *Generalized Conditionalization*, my posterior doxastic attitude should be represented by the interval $[0, 1]$. Below, I explain why.

Suppose W is the space of possibilities over which I distribute my credences. We can partition W into twenty-two cells H_0, H_1, \dots, H_{10} and T_1, \dots, T_{10} . For any i between 0 and 10 (inclusive), any H_i is the set of worlds where a ball is drawn from urn H and the number of blue balls in that urn is i . By contrast, for any i between 0 and 10 (inclusive), T_i is the set of worlds where a ball is drawn from urn T and the number of blue balls in that urn is i . Since a ball will be drawn from urn H if and only if the coin landed heads, *Heads* is simply the union of all the H_i 's: $\bigcup_{i=0}^{10} H_i$. Similarly, since a ball will be drawn from urn T if and only if the coin landed tails, *Tails* is the union of all the T_i 's: $\bigcup_{i=0}^{10} T_i$.

Suppose my rational representor, before I learn anything, is Γ . We know four facts about it.

Fact 1. I have no idea about what the actual distribution of blue and green balls is across the two urns. So, for any i between 0 and 10 (inclusive), $\Gamma(H_i) = \Gamma(T_i) = [0, 1]$.

Fact 2. I am rationally certain that the coin is fair. So, for any probability function p in Γ , $p(\text{Heads}) = p(\text{Tails}) = 0.5$.

Fact 3. The objective chance that a blue ball will be picked from an urn, given that it contains contains i blue balls, should be $\frac{i}{10}$. So, by the Principal Principle, for any p in Γ and any i between 1 and 10 (inclusive), $p(B|H_i) = p(B|T_i) = \frac{i}{10}$.

Fact 4. I am rationally certain that there are exactly ten blue balls in the two urns put together, and that the probability of any particular distribution of balls across the two urns is independent of the outcome of the coin flip. So, any p in Γ , $p(H_i) = p(T_{10-i})$.

Let's see what follows from these four facts. First, consider my prior doxastic state. For

p in Γ ,

$$\begin{aligned}
p(B) &= p(B|Heads)p(Heads) + p(B|Tails)p(Tails) \\
&= \frac{1}{2} \left[\sum_{i=0}^{10} p(B|H_i)p(H_i|Heads) + \sum_{i=0}^{10} p(B|T_{10-i})p(T_{10-i}|Tails) \right] \\
&= \frac{1}{2} \left[\sum_{i=0}^{10} \frac{i \cdot p(H_i|Heads)}{10} + \sum_{i=0}^{10} \frac{(10-i) \cdot p(T_{10-i}|Tails)}{10} \right] \\
&= \frac{1}{2} \left[\sum_{i=0}^{10} \frac{i \cdot p(H_i)}{10p(Heads)} + \sum_{i=0}^{10} \frac{(10-i) \cdot p(T_{10-i})}{10p(Tails)} \right] \\
&= \frac{1}{2} \sum_{i=0}^{10} \frac{(i+10-i)p(H_i)}{10p(Heads)} \\
&= \frac{1}{2} \sum_{i=0}^{10} \frac{p(H_i)}{p(Heads)} \\
&= \frac{1}{2}.
\end{aligned}$$

This means that every probability function in Γ assigns a credence of 0.5 to B . Thus, my prior doxastic attitude towards B should be precise.

Next, consider my posterior doxastic state. Since I update according to *Generalized Conditionalization*, my posterior doxastic attitude towards B in *Heads*-worlds should be represented by $\Gamma(B|Heads) = \{p(B|Heads) : p \in \Gamma\}$, while my posterior doxastic attitude towards B in *Tails*-worlds should be represented by $\Gamma(B|Tails) = \{p(B|Tails) : p \in \Gamma\}$. But note that, for any p in Γ ,

$$\begin{aligned}
p(B|Heads) &= \sum_{i=0}^{10} p(B|H_i)p(H_i|Heads) \\
&= \sum_{i=0}^{10} \frac{i}{10} p(H_i|Heads)
\end{aligned}$$

From *Fact 1* and the fact that the coin flip is independent of the distribution of balls in the two urns, it follows that, for any i , $p(H_i|Heads)$ can be any real number between 0 and 1. So, $p(B|Heads)$ can lie anywhere between 0 and 1. The same is true of $p(B|Tails)$. Thus, my posterior doxastic attitude towards B should be represented by the interval $[0,1]$ irrespective of what I learn. Therefore, if I update my beliefs in this case according to *Generalized Conditionalization*, my doxastic attitude towards B will go from being precise to being imprecise. So, this is a case of dilation.

In this scenario, suppose I am rationally certain before my inquiry that I will update my beliefs according to a conditionalizing rule. So, I can be rationally certain that every probability function in my prior representor assigns a credence of 0.5 to B , and that the minimum credence that B will receive from the posterior representor recommended by my updating rule is 0. In other words, I can be rationally certain that, for

any credence that B might get from a member of my prior representor, a member of the posterior representor recommended by my updating rule will assign a lower credence to B . Therefore, in this case, the conditionalizing rule will be unreflective. I will show in the next four sections (§§5-8) that this failure of *Generalized Weak Reflection* leads to failures of *Value of Evidence*.

5 Good's Theorem

I will begin with Good's (1967) theorem. It is significant, because it entails a restricted version of *Value of Evidence* along with other assumptions.

To state Good's theorem, I need two notions: the notion of a *decision problem* and the notion of an *decision rule*. A *decision problem* is a situation where an agent has to choose amongst a number of available acts. We can represent any such situation with a triple $\langle W, A, v \rangle$, where W is a finite possibility space, A is a set of available acts, and v is a value function that maps any act a in A and any world w in W to a real number which reflects the value of performing a in w by lights of the agent. When an agent makes a choice in relation to a decision problem, they may comply with an *decision rule*, a rule that tells them how to act relative to any decision problem in light of any doxastic state. Such a rule can be represented using a function R that maps a decision problem $\langle W, A, v \rangle$ and a representor Γ to an act a in A .

Such decision rules will be subject to constraints of instrumental rationality. For example, suppose we take expected value maximization to be the norm of instrumental rationality for agents with precise doxastic states. Then, we should think that an agent should only comply with a decision rule which, relative to any decision problem $\langle W, A, v \rangle$ and any rational representor $\{p\}$, recommends acts that maximize expected value relative to the probability function p . Here, the expected value of an act a in A relative to p —written as $Exp_p(a)$ —is a weighted average of a 's values in different worlds, where the weights are probabilities assigned to those worlds by p . That is,

$$Exp_p(a) = \sum_{w \in W} p(w)v(a, w).^{21}$$

More generally, we will call a decision rule *EV-maximizing* just in case, relative to any decision problem and any rational representor, the rule always recommends an act that maximizes expected value relative to some probability function in that representor.

²¹I will assume throughout this essay that, in any decision problem, the states of the world don't depend (either epistemically or causally) on the acts that an agent deliberates about. In fact, complications arise when we relax this assumption and accept either evidential or causal decision theory: evidential decision theorists are led to reject *Value of Evidence* in Newcomb-style cases where states of the world epistemically depend on the acts that an agent deliberates about, while causal decision theorists have to reject it in cases where states of the world causally depend on the acts the agent deliberates about. See Skyrms (1990) for discussion of a Newcomb-style case where the evidential decision theorist must reject *Value of Evidence*, and Rabinowicz (2009) and Ahmed (2014, §7.4.1) for an example about buying a suit of armour, where gathering and using cost-free evidence is suboptimal according to causal decision theory.

Thus, if R is an EV-maximizing decision rule, then, for any rational representor Γ , there is a probability function p in Γ such that, for any decision problem $D = \langle W, A, v \rangle$, $R(D, \Gamma) = a$ for any act a in A just in case there is no other act b in A such that $Exp_p(b) > Exp_p(a)$.

Good (1967) proved the following claim.²² Take any partitional inquiry $\langle W, E, \Gamma \rangle$ involving an agent with a precise doxastic state. Suppose that the agent is rationally certain before the inquiry that they will update according to a conditionalizing rule U and will act according to an EV-maximizing rule R throughout. If p is the sole probability function in Γ , then, relative to p , the expected value of acting in light of the posterior doxastic state recommended by U isn't lower than the expected value of acting in light of the prior doxastic state. In other words, the following inequality holds for any decision problem $D = \langle W, A, v \rangle$.

Good's Inequality. For some probability function p in Γ ,

$$\sum_{w \in W} p(w)v(R(D, \Gamma), w) \leq \sum_{w \in W} p(w)v(R(D, U(w)), w).$$

This supports a restricted version of *Value of Evidence*.

Assume that *Generalized Conditionalization* is true, and that expected value maximization is norm of instrumental rationality for agents with precise doxastic states. Then, Good's theorem entails the following claim.

Restricted Value of Evidence. Suppose a piece of evidence is available to an agent with a precise doxastic state for gathering and use at a negligible cost through a partitional inquiry, and the agent is rationally and correctly certain that they will respond to that evidence in an epistemically and instrumentally rational manner. Then, it is instrumentally rational for that agent to gather that evidence and use it for making decisions.

This claim is more restricted than *Value of Evidence* in two ways: it applies only to cases where (i) the agent's inquiry is partitional, and (ii) the agent's prior doxastic state is precise.²³ Below, I will explore whether we can preserve a similarly restricted version of *Value of Evidence* without the second assumption.

6 Three Norms of Instrumental Rationality

In the case of agents with precise doxastic states, expected value maximization may indeed be the norm of instrumental rationality. But it's not obvious which norms of

²²For a more precise statement of the theorem and its proof, see Appendix A.

²³For the observation that failures of partitionality lead to failures of Good's inequality, see Ahmed and Salow (forthcoming), Dorst (forthcoming), and Das (forthco2020). See the earlier footnote 3 for references to discussions of how Good's inequality can fail in cases of dilation. Buchak (2010) and Campbell-Moore and Salow (2020) have also shown that, if a risk-averse agent doesn't comply with an EV-maximizing rule, then Good's inequality can fail.

instrumental rationality apply to agents with imprecise doxastic states. I will focus on three norms of instrumental rationality that are sometimes taken to be suitable for imprecise agents: *Γ-Maximin*, *E-Admissibility*, and *Maximality*.²⁴ All these constraints are *synchronic* in character: according to them, what an agent is permitted to do at a time doesn't directly depend on what they do (or are disposed to do) at other times.

Start with *Γ-Maximin*. Think of an agent's representor as a credal committee, where each member has a different opinion from the others on some matter. When an agent makes a choice in a decision problem, a member of their credal committee may complain that there is another available action whose worst expected value (by lights of some committee member) is better than the worst expected value of the action chosen by the agent. *Γ-Maximin* says that an agent should eliminate such complaints. Thus, *Γ-Maximin* requires an agent to choose an act that maximizes minimum expected value. This implies the following constraint on decision rules.

Γ-Maximin. An agent is permitted by instrumental rationality to comply with a decision rule R only if R maximizes minimum expected value. That is, for any set of probability functions Γ and any decision problem $D = \langle W, A, v \rangle$, and for any act a in A , if $R(D, \Gamma) = a$, then there is no other act b in A such that $\inf\{Exp_p(b) : p \in \Gamma\} > \inf\{Exp_p(a) : p \in \Gamma\}$.

E-Admissibility gives different advice. It says that, when faced with any decision problem, the credal committee should delegate the choice to one of its members, which can then pick an action that is optimal by its own lights. Thus, according to *E-Admissibility*, an agent is rationally permitted to perform a certain action relative to a decision problem only if there is some probability function in their representor, relative to which this action maximizes expected value. This yields the following constraint on decision rules.

E-Admissibility. An agent is permitted by instrumental rationality to comply with a decision rule R only if R is EV-maximizing. That is, for any set of probability functions Γ , there is a probability function p in Γ such that, for any decision problem $\langle W, A, v \rangle$ and any act a in A , if $R(D, \Gamma) = a$, there is no other act b in A such that $Exp_p(b) > Exp_p(a)$.

The third constraint—*Maximality*—says something even weaker: namely, that if a certain option b has greater expected value than another option a for every member of an agent's credal committee, then the agent shouldn't choose a over b . In other words, an agent shouldn't choose an option which is strictly dominated by another option with respect to expected value. We can state the constraint as follows.

²⁴Gilboa and Schmeidler (1989) prove a representation theorem for *Γ-Maximin* while Berger (1985, §4.7.6) and Grünwald and Dawid (2004) defend it. For discussion of *E-Admissibility*, see Levi (1974, 1980, 1986). For discussion of *Maximality*, see Walley (1991). While discussing *Value of Evidence*, Kadane, Schervish and Seidenfeld (2008) examine all three norms, while Bradley and Steele (2016) focus on *Γ-Maximin* and *Maximality*.

Maximality. An agent is permitted by instrumental rationality to comply with a decision rule R only if R recommends acts that are strictly undominated with respect to their expected value. That is, for any set of probability functions Γ , any decision problem $\langle W, A, v \rangle$, and any act a in A , if $R(D, \Gamma) = a$, there is no other act b in A such that, for any p in Γ , $Exp_p(b) > Exp_p(a)$.

If an agent complies with a decision rule that satisfies either Γ -*Maximin* or *E-Admissibility*, they will also comply with a decision rule that satisfies *Maximality*. If an act were strictly dominated by another with respect to expected value, then it couldn't maximize minimum expected value. For there would be another act that has higher minimum expected value than it. Similarly, if an act were strictly dominated by another act with respect to expected value, then it couldn't maximize expected value relative to any probability function in the agent's representor. For there would be another act that has higher expected value than it relative to any probability function in the agent's representor.

For agents with rational precise doxastic states, Γ -*Maximin*, *E-Admissibility*, and *Maximality* yield the same predictions: they require the agent to conform to an EV-maximizing decision rule. But they diverge in other cases. Take a decision problem, involving three options a , b , and c , whose payoffs are given in Table 1.²⁵

	X	$\sim X$
a	1	0
b	0	1
c	0.4	0.4

Table1: A Case where Γ -*Maximin* and *E-Admissibility* Diverge

Suppose my doxastic state is imprecise, so I assign an interval of credences $[0.1, 0.9]$ to X . If a decision rule R satisfies Γ -*Maximin*, then, relative to my prior doxastic state, R will recommend option c . This is because the set of expected values for both a and b is $[0.1, 0.9]$ relative to my representor, and the set of expected values for c is $\{0.4\}$. So, c maximizes minimum expected value. By contrast, if R satisfies *E-Admissibility*, then R will recommend either a or b , but definitely not c . This is because, relative to any probability function in Γ , either a or b maximizes expected value, since the expected value of at least one of these options is always greater than or equal to 0.5. Thus, decision rules that satisfy Γ -*Maximin* and *E-Admissibility* will recommend different options in this case. Finally, if R satisfies *Maximality*, then R could recommend all three options, since none of the options are strictly dominated by the others.

²⁵This example is due to Seidenfeld (2004).

7 A Failure of Good's Inequality

I want to show that, in cases where *Generalized Weak Reflection* fails, decision rules that satisfy these constraints can lead to failures of a weakly restricted version of *Value of Evidence*:

Weakly Restricted Value of Evidence. Suppose a piece of evidence is available to an agent for gathering and use at a negligible cost through a partitioned inquiry, and the agent is rationally and correctly certain that they will respond to that evidence in an epistemically and instrumentally rational manner. Then, it is instrumentally rational for that agent to gather that evidence and use it for making decisions.

The difference between *Weakly Restricted Value of Evidence* and *Restricted Value of Evidence* is that the latter only applies to agents with precise doxastic states, but the former doesn't. To show how *Weakly Restricted Evidence* can be false, I will focus on a version of *Mystery Urn 2* where I am rationally certain before my inquiry about which decision rule I will comply with.

In *Mystery Urn 2*, every member of my rational prior representor assigns a precise credence of 0.5 to the proposition that a blue ball will be chosen. But, after I update my beliefs according to *Generalized Conditionalization*, I can only rationally assign an interval of credences $[0, 1]$ to that proposition. Suppose that, before I engage in this inquiry, I am offered a bet. If I accept it, I gain 2 units of utility when a blue ball is chosen but lose 1 unit of utility when a green ball is chosen. If I reject the bet, I get nothing. Thus, the relevant decision problem involves two options: *Accept* (the option of accepting the bet) or *Reject* (the option of rejecting the bet). The payoffs for these options are given in Table 2. Suppose I can either make a decision about this bet before

	B	$\sim B$
<i>Accept</i>	2	-1
<i>Reject</i>	0	0

Table2: A Decision Problem for *Mystery Urn 2*

my inquiry, or after. So, I have the option of delaying my decision until after my inquiry. Should I take it?

Let R be any decision rule that satisfies Γ -maximin. For any member of my rational prior representor, the expected value of *Accept* is $0.5 \times 2 + 0.5 \times -1 = 0.5$ and the expected value of *Reject* is 0. So, R will recommend that I accept the bet in light of my prior doxastic state. Things are different relative to my posterior doxastic state. Since my posterior doxastic attitude towards B should be represented by the interval $[0, 1]$, the set of expected values of *Accept* is $[-1, 2]$ while the set of expected values of *Reject* is $\{0\}$. So, the minimum expected value of *Reject* is 0, while the minimum expected value

of *Accept* is -1. Therefore, *R* will recommend that I reject the bet in light of my posterior doxastic state.

Suppose *Generalized Conditionalization* is right, and Γ -maximin is a norm of instrumental rationality. And suppose that, before my inquiry, I am rationally and correctly certain that I will rationally revise my beliefs according to a conditionalizing rule and that I will rationally act according to *R*. So, I can be rationally certain of two things: (i) that, relative to my prior doxastic state, I will accept the bet, and (ii) that, relative to my posterior doxastic state, I will reject the bet no matter what I learn. So, relative to every probability function in my rational prior representor, the expected value of acting in light of my prior doxastic state will be the same as that of *Accept*: 0.5. By contrast, the expected value of acting in light of my posterior doxastic state will be the same as that of *Reject*: 0. Therefore, *Good's Inequality* will fail. So, by *Maximality*, it will be irrational for me to gather the available evidence and use it for making a decision about the bet even though the evidence is cost-free and I am rationally certain of my own rationality. Thus, *Weakly Restricted Value of Evidence* will be false.

Something similar, but weaker, is true of *E-Admissibility* and *Maximality*. Notice that the Γ -*Maximin*-satisfying decision rule *R* is EV-maximizing: it maximizes expected value relative to some probability function in my representor at every stage. So, it satisfies both *E-Admissibility* and *Maximality*. Therefore, in *Mystery Urn 2*, there is a decision rule *R* that satisfies *E-Admissibility* and *Maximality* such that, if I am antecedently rationally certain that I will comply with *R*, *Good's Inequality* will fail. If *Generalized Conditionalization* is right and *R* is rational for me to comply with it, then it can be instrumentally irrational for me in this case to gather and use cost-free evidence before making my decision, even though I may be rationally and correctly certain that I will rationally respond to my evidence throughout. Thus, *E-Admissibility* and *Maximality* will allow *Weakly Restricted Value of Evidence* to be false.

8 A Diagnosis

We can generalize our observations from the previous section.

Proposition 1. For any inquiry $\langle W, E, \Gamma \rangle$, suppose the relevant agent is rationally certain that they will update according to an unreflective rule *U*. Then, the following claims are true.

- *Claim 1.1.* For any decision rule *R* that satisfies Γ -*Maximin*, *Good's inequality* fails for *R* and some decision problem *D*.
- *Claim 1.2.* There exists a decision rule *R* that satisfies *E-Admissibility* and *Maximality* such that *Good's inequality* fails for *R* and some decision problem *D*.²⁶

²⁶All proofs are in Appendix B.

Let's flesh out the significance of the two claims. *Claim 1.1* implies that, if an agent is rationally certain before their inquiry that they will update according to an unreflective rule and that they will comply with Γ -maximin, there is a decision problem relative to which Good's inequality will fail. Thus, *Claim 1.1* confirms a point that Kadane, Schervish and Seidenfeld (2008) and Bradley and Steele (2016) have already made. It shows us that, given *Generalized Conditionalization* and Γ -Maximin, *Rationality of Imprecision* can conflict with *Value of Evidence*. Suppose *Rationality of Imprecision*, *Generalized Conditionalization*, and Γ -Maximin are true. Then, an agent can indeed be correctly and rationally certain that they will rationally revise their beliefs using an unreflective conditionalizing rule and that they will rationally act according to a decision rule that satisfies Γ -Maximin. Then, according to every member of their prior rational representor, the option of acting in light of their prior doxastic state will have greater expected value than the option of acting in light of their posterior doxastic state. As a result, they will be rationally required not to gather and use a piece of available evidence even if it is cost-free. Thus, *Weakly Restricted Value of Evidence* will be false.

Kadane, Schervish and Seidenfeld (2008) and Bradley and Steele (2016) are more optimistic about *E-Admissibility* and *Maximality* respectively: they think that *E-Admissibility* or *Maximality* can help us reconcile *Rationality of Imprecision* with *Value of Evidence*. This is because they focus on agents who have no clue about what their future choices in different possible situations will be.²⁷ By contrast, I have focused on agents who are rationally certain about which decision rule they will comply with. There is nothing irrational about such an agent: in a situation where there are multiple decision rules that an agent is rationally permitted to use, the agent may indeed arbitrarily commit themselves to one and thereby become rationally certain that they will comply with that decision rule.

Claim 1.2 shows that Good's inequality may fail for such an agent. It implies that, if such an agent is rationally certain that they will revise their beliefs according to an unreflective updating rule and act according to a *certain* decision rule that satis-

²⁷See, especially, Bradley and Steele (2016, §3.2). Their strategy (given in §4 of their paper) for defending something like *Weakly Restricted Value of Evidence* is roughly this. In a case where the agent has no clue about what their future choices will be, they at best will be able to single out a set of permissible acts that they could perform when they receive a certain piece of evidence. For example, this could be the set of *E-admissible* or *maximal* options. Since the agent cannot pin down exactly how they will act, they won't be able to assign a precise expected value to the option of acting in light of their future doxastic state. Rather, they will assign a set of expected values. Bradley and Steele (2014) then appeal to a principle, which is more commonly known as *Interval Dominance* but they call the "non-dominated set" rule. The basic idea is that the agent can strictly prefer to act in light of their prior doxastic state only if each number in the set of expected values of that option is higher than each number in the set of expected values of acting in light of their future doxastic state. This cannot happen given Good's theorem, because, for each probability function p in the agent's representor, the expected value of performing acts that maximize expected value relative to conditionalized versions of p will at least be as high as the expected value of performing acts that maximize expected value relative to p . So, Bradley and Steele claim that *E-Admissibility* or *Maximality* can preserve *Weakly Restricted Value of Evidence*. I am trying to argue that this isn't quite right, since they have overlooked the case of agents who are antecedently rationally certain about which decision rule they will comply with. For such agents, at least, the possibility that Good's inequality might fail is left open.

fies *E-Admissibility* and *Maximality*, then there will be a decision problem relative to which Good's inequality will fail. So, if *Generalized Conditionalization* is true and *E-Admissibility* or *Maximality* are norms of instrumental rationality, then there *may* be a tension amongst *Rationality of Imprecision* and *Value of Evidence*. If *Rationality of Imprecision* and *Generalized Conditionalization* are true, an agent can rationally engage in an unreflective inquiry. Assume that it is also instrumentally rational for that agent to comply with an *E-Admissibility*- or *Maximality*-satisfying decision rule that violates Good's inequality in such an unreflective inquiry. Then, if the agent is rationally and correctly certain that they will rationally revise their beliefs using an unreflective updating rule and will rationally comply with the relevant decision rule, then they will be rationally required not to gather cost-free evidence and use it for making decisions. So, *Weakly Restricted Value of Evidence* will be false.

Can defenders of *E-Admissibility* and *Maximality* block this argument? I think they can. They can simply reject the assumption that it can be instrumentally rational for an agent to comply with an *E-Admissibility*- or *Maximality*-satisfying decision rule that violates Good's inequality in an unreflective inquiry. If that assumption is false, it will be instrumentally irrational for an agent to comply with such a decision rule. So, an agent who is rationally and correctly certain of their own rationality can't also be rationally certain that they will comply with such a rule. For, if they were rationally certain that they would comply with such a decision rule, then their certainty regarding their own rationality would be misplaced; they would be wrong to be certain that they would respond to their evidence in an instrumentally rational manner. But no one—as far as I know—has explored this strategy. In §§10-11, I do this.

9 Consequences of the Diagnosis

Before I do this, I want to address an objection. One might worry that *Proposition 1* is too weak: it cannot rule out the possibility that there are other norms of instrumental rationality that can reconcile *Rationality of Imprecision* with *Weakly Restricted Value of Evidence*.

I agree. What *Proposition 1* shows is that there is no obvious *synchronic* norm of instrumental rationality that can help us resolve the tension between *Rationality of Imprecision* and *Weakly Restricted Value of Evidence*. We can begin to see this simply by considering three other synchronic norms that are either stronger than or compatible with *E-Admissibility* and *Maximality*.

The first norm I want to consider is *Levi's Rule*.²⁸ It combines *E-Admissibility* with *Γ -Maximin*. It says that an agent should choose an option *a* just in case (a) *a* is *E*-admissible (i.e., maximizes expected value relative to some probability function in the agent's representor), and (b) *a* maximizes minimum expected value amongst all the

²⁸It is defended by Levi (1974, 1980, 1986).

E-admissible options. We can state this rule more precisely as follows.

Levi's Rule. An agent is permitted by instrumental rationality to comply with a decision rule R iff R recommends an E-admissible act that maximizes minimum expected value amongst E-admissible acts. That is, for any set of probability functions Γ , any decision problem $D = \langle W, A, v \rangle$ and any act a in A , $R(D, \Gamma) = a$ iff two conditions are satisfied:

- R is EV-maximizing.
- If A^* is the set of E-admissible options in A , there is a Γ -maximin-satisfying decision rule R^* such that $R^*(\langle W, A^*, v \rangle, \Gamma) = a$.

This norm is stronger than *E-Admissibility*. But, insofar as it incorporates a version of Γ -*Maximin*, it inherits the problems of Γ -*Maximin*. For example, in *Mystery Urn 2*, if I am offered the bet that pays 2 if B is true and -1 otherwise, both accepting the bet and rejecting it are E-admissible relative to my future doxastic state (since both maximize expected value relative to some member of my posterior representor). But the set of expected values of accepting the bet is $[-1, 2]$ while the set of expected values of rejecting the bet is $\{0\}$. So, only rejecting the bet maximizes minimum expected value. But, as we know, if I am antecedently rationally certain that I will reject the bet in light of my posterior doxastic state in this case, then Good's inequality will fail. Thus, Levi's rule cannot help us preserve Good's inequality. We can in fact derive the following corollary from *Proposition 1*.

Corollary 1. For any partitional inquiry $\langle W, E, \Gamma \rangle$, suppose the relevant agent is rationally certain that they will update according to an unreflective rule U . Then, if a decision rule R satisfies *Levi's Rule*, then Good's inequality fails for R and some decision problem D .

So, we cannot hope to reconcile *Weakly Restricted Value of Evidence* with *Rationality of Imprecision* by appealing to *Levi's Rule*.

One might think that the culprit here is Γ -maximin. So, consider another norm called *Caprice*—proposed by Weatherson (1998) and discussed by Williams (2014) and Moss (2015)—which doesn't involve any application of Γ -maximin. According to this norm, an agent is rationally permitted to perform an act if and only if some members of their credal committee regard it as optimal. This suggests a relatively simple constraint on decision rules.

Caprice. An agent is permitted by instrumental rationality to comply with a decision rule R iff R is EV-maximizing.

Like *Levi's Rule*, this norm is stronger than *E-Admissibility*. But *Claim 1.2* in *Proposition 1* shows why this can't preserve *Weakly Restricted Value of Evidence*. In *Mystery Urn 2*, the decision rule R that I comply with is an EV-maximizing rule. So, if *Caprice* is right,

then it can be instrumentally irrational for an agent who is rationally certain of their own rationality to gather and use cost-free evidence.

Finally, take the norm—proposed originally by Hurwicz (1951)—called the *Hurwicz Criterion*. This norm can be stated in terms of the *Hurwicz expected value* of an act, which is just a weighted average of its minimum expected value and its maximum expected value according to probability functions in the agent’s representor. In other words, for any decision problem $\langle W, A, v \rangle$ and any set of probability functions Γ , if a is an act in A , then its Hurwicz expected value relative to Γ is given by:

$$HExp_{\Gamma, \alpha}(a) = \alpha \cdot (\inf\{Exp_p(a) : p \in \Gamma\}) + (1 - \alpha) \cdot (\sup\{Exp_p(a) : p \in \Gamma\}),$$

for some α between 0 and 1 (inclusive). The *Hurwicz Criterion* says that an agent should act just in case it maximizes Hurwicz expected value. This yields the following constraint.

The Hurwicz Criterion. An agent is permitted by instrumental rationality to comply with a decision rule R iff R maximizes Hurwicz expected value. That is, for any set of probability functions Γ , any decision problem $D = \langle W, A, v \rangle$, and any act a in A , $R(D, \Gamma) = a$ iff there is no other act b in A such that $HExp_{\Gamma, \alpha}(b) > HExp_{\Gamma, \alpha}(a)$ for some α between 0 and 1 (inclusive).

In some cases where an agent’s rational representor is convex, the *Hurwicz Criterion* can be compatible with *E-Admissibility* (given certain values of α). In those cases, it only recommends acts that maximize expected value according to some probability function in the agent’s representor. For example, in *Mystery Urn 2*, the set of expected values of accepting the bet relative to my prior doxastic state is $\{0.5\}$. So, its Hurwicz expected value is 0.5, which is greater than than Hurwicz expected value of rejecting the bet (i.e., 0). So, a decision rule that satisfies the *Hurwicz Criterion* for any α will recommend that I act the bet relative to my prior doxastic state. By contrast, the set of expected values of accepting the bet relative to my posterior doxastic state is $[-1, 2]$. So, for any α between 0 and 1 (inclusive), the Hurwicz expected value of accepting the bet is $\alpha \cdot (-1) + (1 - \alpha) \cdot 2 = 2 - 3\alpha$. If α is greater than $\frac{2}{3}$, the Hurwicz expected value of accepting the bet will be negative. So, a decision rule that satisfies the *Hurwicz Criterion* for such an α will recommend that I reject the bet relative to my posterior doxastic state. This rule is exactly like the EV-maximizing rule that I described in the last section. So, the problem remains the same: if I am rationally certain in advance that I shall comply with such a rule, then Good’s inequality fail. This observation can be generalized.

Corollary 2. For any partitional inquiry $\langle W, E, \Gamma \rangle$, suppose the relevant agent is rationally certain that they will update according to an unreflective rule U . Then, there is a decision rule R that satisfies the *Hurwicz Criterion* for some α between 0.5 and 1 (exclusive), such that Good’s inequality fails for R and some decision problem D .

This shows that the *Hurwicz Criterion* can't preserve Good's inequality.²⁹

All these norms are synchronic in character: according to them, what an agent is permitted to do at a time doesn't directly depend on how they act (or are disposed to act) at other times. In the next section, I show that there is a diachronic norm of instrumental rationality which escapes the problems that arise for these synchronic norms.

10 A Diachronic Norm of Rationality

The diachronic norm I have in mind is *Practical Stability*.

Call a decision rule *stable* just in case it recommends actions that maximize expected value relative to probability functions that diachronically cohere with each other in the course of an inquiry. a decision rule is stable relative to a probability function p in the agent's rational prior representator—in other words, p -stable—just in case it satisfies two conditions. First, prior to the inquiry, it recommends an action that maximizes expected value relative to p . Second, given evidence E_i , it recommends an action that maximizes expected value relative to a probability function $p(.|E_i)$. As usual, we can state the idea more formally.

Stable Decision Rules. For any inquiry $\langle W, E, \Gamma \rangle$ and for any probability function p in Γ , a decision rule R is p -stable iff:

- For any decision problem $D = \langle W, A, v \rangle$ and any act b in A , $R(D, \Gamma) = b$ iff $b \in \arg \max_{a \in A} Exp_p(a)$.
- For any decision problem $D = \langle W, A, v \rangle$, any world w in W and any act b in A , $R(D, U(w)) = b$ iff $b \in \arg \max_{a \in A} Exp_{p(.|E(w))}(a)$.

It's easy to see that the recommendations of a stable decision rule will diachronically cohere in a certain way. For any probability function p in the agent's rational prior representor, if a decision rule is p -stable, then it will not only recommend an action that

²⁹This point also extends to a similar decision rule proposed by Ellsberg (1961, pp. 664). Ellsberg's rule appeals to two notions: the *estimated value* of an act and the *index value* of an act. The estimated value of an act is simply its expected value calculated according to a weighted average of the probability functions in the relevant agent's rational representor. Let this be $Est(a)$. The index value of an action a is the weighted average of the minimum expected value of a and the estimated value of a : $\alpha \cdot \inf\{Exp_p(a) : p \in \Gamma\} + (1 - \alpha) \cdot Est(a)$, where α lies between 0 and 1 (inclusive). According to Ellsberg's rule, an agent should perform acts that maximize this index value. In *Mystery Coin 2*, before I gather my evidence, since *Accept* maximizes expected value according to every member of my rational representor, that option will maximize the index value. But, after I gather my evidence, if the estimated value is calculated according to the weighted average of my posterior representor (such that I assign equal weight to each member of my representor), the estimated value of *Accept* will be 0.5. But, as we know, the minimum expected value of *Accept* is -1. For any α between 0 and 1, then the index value of *Accept* will be $0.5 - 1.5\alpha$, which is negative if α is greater than $\frac{1}{3}$. In that case, *Reject* will maximize the index value of the act. So, Ellsberg's rule will recommend that I reject the bet after I gain new evidence. Suppose I comply with Ellsberg's rule. Then, since I can be rationally certain in this case that I initially accept the bet but will later reject it, Good's inequality will fail.

maximizes expected value according to p relative to the agent's prior doxastic state, but also, given evidence E_i , it won't recommend an action that fails to maximize expected value relative to $p(\cdot|E_i)$. We can now show that Good's inequality fails for the decision rules we've discussed so far because they are not stable in the relevant sense.

In *Mystery Urn 2*, suppose that I am rationally certain before the inquiry that I will comply with a p -stable decision rule R . So, there is some probability function p in my representor, such that (i) relative to my prior doxastic state, R recommends an action that maximizes expected value relative to p ; (ii) when I learn *Heads*, R recommends an action that maximizes expected value relative to $p(\cdot|Heads)$; and (iii) when I learn *Tails*, R recommends an action that maximizes expected value relative to $p(\cdot|Tails)$.

First, consider what R recommends relative to my prior representor when I face the bet that pays 2 if B is true and -1 otherwise. Since accepting the bet uniquely maximizes expected value relative to each member of my representor, R will recommend that I accept the bet. Next, consider what R recommends relative to my posterior representor. If R recommends that I reject the bet when I learn *Heads*, this can only be because, relative to $p(\cdot|Heads)$, rejecting the bet maximizes expected value. And this can only happen if $p(B|Heads) \leq \frac{1}{3}$. But, then, $p(B|Tails) \geq \frac{2}{3}$.³⁰ Thus, relative to $p(\cdot|Tails)$, rejecting the bet won't maximize expected value. So, since R is p -stable, it won't recommend that I reject the bet when I learn *Tails* (provided that it recommends that I reject the bet when I learn *Heads*). Thus, in this case, I cannot be rationally certain that I will reject the bet relative to my posterior doxastic state regardless of what I learn.

This can help us preserve Good's inequality. For example, suppose $p(B|Heads) \leq \frac{1}{3}$ and $p(B|Tails) \geq \frac{2}{3}$. So, let R recommend that I turn down the bet when I learn *Heads* but accept it when I learn *Tails*. According to p , the expected value of acting in light of my future doxastic state will be at least be $0.5 \times 0 + 0.5 \times (\frac{2}{3} \times 2 + \frac{1}{3} \times (-1)) = 0.5$. Thus, the expected value of acting in light of my future doxastic state will be at least as great as the expected value of acting in light of my prior doxastic state. So, Good's inequality won't fail.

The norm of rationality that I wish to endorse is this.

Practical Stability. In the context of any partitional inquiry $\langle W, E, \Gamma \rangle$, if R is a decision rule that an agent is able to comply with, then they are permitted by instrumental rationality to comply with R iff R is p -stable for some probability function p in Γ .

³⁰This is because:

$$\begin{aligned} p(B) &= p(Heads)p(B|Heads) + p(Tails)p(B|Tails) \\ \Rightarrow 0.5 &= 0.5(p(B|Heads) + p(B|Tails)) \\ \Rightarrow p(\sim B|Heads) &= p(B|Tails) \end{aligned}$$

If $p(B|Heads) \leq \frac{1}{3}$, then $p(\sim B|Heads) \geq \frac{2}{3}$. So, $p(B|Tails) \geq \frac{2}{3}$.

This is a diachronic norm of rationality: according to this norm, what an agent is permitted to do at a certain stage of an inquiry depends on what they do (or are disposed to do) at other stages of that inquiry. Note what *Practical Stability* doesn't say. First, it doesn't say that an agent is *always* required to comply with stable decision rules. That would be too strong: indeed, in a case where an agent has no control over what their future self would do, complying with a stable decision rule may not be an available option. *Practical Stability* includes a proviso to rule out such cases: it only applies to decision rules that the agent is able to comply with. Second, it doesn't say that an agent who is able to comply with a stable decision rule across different inquiries must comply with the same stable decision rules across different inquiries. That would collapse the distinction between a rational agent with a precise doxastic state and a rational agent with an imprecise doxastic state, since both would be performing actions that maximize expected value relative to conditionalized versions of the same prior probability function. Rather, *Practical Stability* says that, *in the context of a partitional inquiry*, an agent is permitted by rationality only to comply with stable decision rules (provided that they are able to do so). So, in different inquiries, the agent may comply with different stable decision rules.

We can show that *Practical Stability* preserves Good's inequality.

Proposition 2. For any partitional inquiry $\langle W, E, \Gamma \rangle$, any p in Γ , any p -stable decision rule R and any decision problem D , Good's inequality will hold relative to p .

Proposition 2 sheds light on a connection between *Practical Stability* and *Value of Evidence*. Suppose *Practical Stability* is true. Suppose that an agent is able to comply with a stable decision rule in the context of a partitional inquiry, and is rationally certain about which decision rule they will comply with. If that agent is rationally and correctly certain of their own rationality, then they will be rationally certain that they will comply with a p -stable decision rule throughout their inquiry (for some p in their rational prior representor). *Proposition 2* shows that Good's inequality will hold for the relevant p . In other words, by lights of p , it will be optimal for the agent gather the available cost-free evidence and use it for making their decision. So, the relevant p -stable decision rule will recommend that the agent do so. Then, by *Practical Stability*, it will be instrumentally rational for the agent to gather the available evidence and use it for making their decision. In this way, *Practical Stability* preserves *Weakly Restricted Value of Evidence* for agents who are able to comply with stable decision rules and are rationally certain about which decision rules they will comply with.

11 Can Practical Stability Be Defended?

I have shown that *Practical Stability* will help us preserve a restricted version of *Value of Evidence*. But I haven't considered whether a defender of *Rationality of Imprecision*

should accept *Practical Stability*.

In this section, I wish to argue that we can partially defend this norm by appealing to a more general principle about instrumental rationality:

The Principle of Exploitability. If R is a decision rule that an agent is able to comply with, then they are permitted by instrumental rationality to comply with R only if complying with R doesn't make them predictably exploitable.³¹

If *E-Admissibility* is a genuine norm of rationality and the *Principle of Exploitability* is right, then *Practical Stability* must be true. This is because complying with an unstable but EV-maximizing rule makes an agent predictably vulnerable to a *diachronic Dutch book*, i.e., a series of bets which are offered at different times and, when accepted, result in a net loss of utility.

To see why this is true, consider *Mystery Urn 2* once more. Suppose—for *reductio*—that, in this case, I am rationally and correctly certain that I will rationally act according to an unstable EV-maximizing rule R such that, irrespective of whether I learn *Heads* or *Tails*, R requires me to maximize expected value according to a probability function p in my posterior representor that assigns 0.1 to B .³² Then, I can be offered a bet a before my inquiry and a bet b after my inquiry, with the payoffs given in Table 3. In my

	B	$\sim B$
a	0.6	-0.4
b	-0.8	0.2

Table3: A Diachronic Dutch book for the Practically Unstable Decision Rule R in *Mystery Urn 2*

prior doxastic state, every probability function in my representor assigns a credence of 0.5 to B . So, for any such probability function, the expected value of accepting a is $0.6 \times (0.5) + 0.5 \times (-0.4) = 0.1$, which is greater than the expected value of rejecting a (i.e., 0). So, if I comply with an EV-maximizing rule, I will accept a . When I learn about the outcome of the coin flip, my doxastic attitude towards B becomes imprecise; it can now be represented by the interval $[0, 1]$. We know that R maximizes expected value relative to a probability function that assigns a credence of 0.1 to B . According to any such probability function, the expected value of accepting b will be $0.1 \times (-0.8) + 0.9 \times 0.2 = 0.1$, which is greater than the expected value of rejecting that bet (i.e., 0). So, R will recommend that I accept that bet. But, if I accept both bets, I am subject to a net loss

³¹Versions of the *Principle of Exploitability* are somewhat popular amongst epistemologists and decision theorists. Proponents of money pump arguments for norms of instrumental rationality and Dutch book arguments for norms of epistemic rationality typically assume that a rational agent cannot be predictably exploitable. See, for example, Ramsey (1926), de Finetti (1992), and Davidson, McKinsey and Suppes (1955).

³²This rule is unstable because, for any p in my prior representor, if $p(B|Heads) = 0.1$, then $p(B|Tails) = 0.9$.

of 0.2. Then, according to the *Principle of Exploitability*, if I am able to comply with R in the course of this inquiry, it cannot be instrumentally rational for me to comply with R . So, I cannot be rationally and correctly certain that I will rationally act according to R . This contradicts our initial assumption.³³

However, if I am able to comply with a stable decision rule in this case, then *Practical Stability* will permit me to act according to such a rule throughout the inquiry. If I comply with such a rule, I won't accept b irrespective of what I learn. To check this, suppose my decision rule is p -stable for some p in my rational prior representor. Since $p(B) = 0.5$, the expected value of accepting a is positive according to p . So, I will initially accept bet a . But if I learn that the coin landed heads, I can only accept b if the expected value of accepting b is non-negative according to $p(\cdot|Heads)$. And that can happen only if $p(B|Heads) \leq \frac{1}{5}$. But then $p(B|Tails) \geq \frac{4}{5}$ (by the reasoning given in footnote 30). In that case, the expected value of accepting b will be negative according to $p(\cdot|Tails)$. So, I must turn down the bet when I learn that the coin landed tails. Similarly, if I learn that the coin landed tails, I can only accept b if $p(B|Tails) \leq \frac{1}{5}$. But, then, $p(B|Heads) \geq \frac{4}{5}$. So, I must turn down the bet when I learn that the coin landed heads. In other words, if I comply with a stable decision rule, I won't necessarily accept bet b after my inquiry is over. As a result, I won't be predictably exploitable.

The point can be generalized. Suppose the strongest pieces of evidence that an agent could get in the course of an inquiry are E_1, E_2, \dots, E_k . Let a decision rule R be subject to a diachronic Dutch book just in case there is a series of bets b, b_1, \dots, b_k such that (i) b is offered to the agent before their inquiry and each b_i is offered just in case the agent receives E_i as their posterior evidence, and (ii) if the agent complies with R relative to her prior and posterior doxastic state, then they will accept all the bets and undergo a net loss of utility. We can show that:

Proposition 3. For any partitional inquiry $\langle W, E, \Gamma \rangle$, suppose the agent is rationally certain that they will update according to a conditionalizing rule U . Then, if R is an EV-maximizing rule, it is not subject to a diachronic Dutch book iff it is p -stable for some p in Γ .

This is significant. If *E-Admissibility* is a genuine norm of instrumental rationality, then an agent who updates according to *Generalized Conditionalization* is permitted only to comply with an EV-maximizing decision rule in the context of a partitional inquiry.

³³An interesting feature of this Dutch book argument is that it is immune to a response that Schick (1986) offer. Schick points out that a *sophisticated chooser*—who sees the Dutch book coming and is rationally certain that they will accept one (or more) of the subsequent bets—will turn down the initial bet(s) of the Dutch book (by engaging in a form of backward induction). So, they won't be exploitable. This strategy fails in this case. Suppose, before my inquiry, I am rationally certain that I will face choices pertaining to both a and b before and after my inquiry, and that I will accept bet b . Given this, I can either accept or reject a at the initial stage. By my lights, the expected value of accepting a (and then b later) is -0.2, since I will lose 0.2 no matter what happens. By contrast, the expected value of rejecting a (and then accepting b later) is $0.5 \times (-0.8) + 0.5 \times 0.2 = -0.3$. So, I will accept a . Thus, I will be predictably exploitable.

But, unless such a decision rule is stable, the agent will be predictably exploitable. So, in such an inquiry, the agent is permitted only to act according to a stable decision rule (provided that they are able to comply with it). This partly supports what *Practical Stability* says.

In the next three sections (§§12-14), I want to address three challenges for *Practical Stability*.

12 Challenge 1: Exploitation

In §11, I motivated *Practical Stability* by appealing to the idea that using stable decision rules helps us avoid being predictably exploitable. But this may not be convincing: using stable decision rules *merely* within the context of an inquiry cannot prevent an agent from being subject to a diachronic Dutch book. This, in turn, may undercut the motivation for *Practical Stability*.³⁴

Recall *Mystery Coin 1* from §1. There, I assigned an interval of credences $[0,1]$ to B , i.e., the proposition that the ball selected from the urn will be blue. Let's extend the story. Suppose there are two fair coins, 1 and 2, that will be flipped in succession. By observing the coin flips, I will first learn whether coin 1 lands heads or tails, and then whether coin 2 lands heads or tails. Since the coin flips are independent of the colour of the ball drawn from the urn, my doxastic attitude towards B will remain unchanged throughout. Assume that these two episodes of learning are distinct partitioned inquiries. Then, *Practical Stability* will allow me pick a member of my representor, p , to make decisions in the context of the first inquiry, and then switch another member of my representor, p^* , at the beginning of the second. In this extended version of *Mystery Coin 1*, there are four possible bodies of evidence I could end up with as a result of the two inquiries (since there are four combinations of outcomes of the two coin flips). Let HH be the proposition that both coins land heads, and let $p(B) \neq p^*(B|HH)$. Then, if I accept bets at the beginning of the first inquiry by using p and at the end of the second inquiry by using $p^*(\cdot|HH)$ after I have learnt HH , I will be subject to a diachronic Dutch book (by the same reasoning that underlies *Proposition 3*). Therefore, even though I conform to *Practical Stability* in this case, it doesn't prevent me from being exploitable. So, the motivation for *Practical Stability*—namely, that adhering to stable decision rules prevents us from being exploitable—turns out to be wrong.

This challenge can be resolved. To do this, we need to distinguish two versions of this example: a version where the two episodes of learning are part of a single partitioned inquiry and a version where they are not.

Begin with the first version of the case. On my conception of an inquiry, in order to count as an inquiry, an evidence-gathering act or event must satisfy two conditions: (i) prior to that act or event, the relevant agent can be rationally certain that they will

³⁴Thanks to an anonymous referee for suggesting this objection.

engage in that act or undergo that event, and (ii) the relevant agent won't lose any evidence in the course of that evidence-gathering act or event (but might gain some evidence). Imagine a version of this case where, before the coin flips, I am or can be rationally certain that I will engage in some evidence-gathering act or event, as part of which I will (or might) observe these coin flips. Suppose also that, in the course of observing the coin flips, I won't lose any evidence I already possess, and that my posterior evidence will both be factive and have perfect introspective access to itself. Then, according to conditions (i) and (ii), the longer evidence-gathering act or event that includes the two observations of the coin flips should also be treated as a partitioned inquiry. Then, *Practical Stability* will say that I should (if I can) comply with a stable decision rule relative to this longer inquiry. If I comply with such a rule, I won't be predictably exploitable.

This shows that the aforementioned problem of exploitability can only arise if these two episodes of learning cannot be embedded within a single partitioned inquiry. And that can happen if (a) I am not or can't be antecedently rationally certain that I will take part in some evidence-gathering act or event, in the course of which I will (or might) observe these two coin flips, or if (b) I lose evidence in the course of making these two observations, or if (c) my posterior evidence is either not factive or lacks perfect introspective access to itself.

Take a case where (a) is true. In such a case, I am not or can't be rationally certain that I will participate in an evidence-gathering act or event in the course of which I will (or might) observe these coin flips. For example, before the first coin flip, I might be rationally certain that I will observe it, without anticipating or even entertaining the possibility that I might observe the second coin flip. But, then, once I've observed the first coin flip, I might learn that I will observe a second coin flip. As a result, the two observations I make might end up being distinct inquiries that cannot be embedded within the same inquiry. If that happens, I might be permitted by *Practical Stability* to use a decision rule that doesn't remain stable across those two inquiries. However, while this might make me *actually* exploitable, I may not be able to foresee that I will undergo a sure loss of utility if my decision rule doesn't remain stable in the course of observing the two coin flips. Thus, I won't be *predictably* exploitable. So, I won't be required by the *Principle of Exploitability* to comply with a decision rule that remains stable across the two inquiries.

Consider a case where (b) is true. This will be a case of evidence loss, e.g., a case where I will forget some piece of evidence X when I gain the new evidence E_i . Suppose that I am epistemically rational. Since I lose the relevant piece of evidence, the set of credences I assign to X in this case should go from being $\{1\}$ to being some interval of credences lower than 1 when I learn E_i .³⁵ As a result, my posterior set of credences

³⁵I am assuming that, when the agent distributes credences over a finite possibility space, the credence functions in the agent's posterior representor are required by epistemic rationality to assign non-zero

in X cannot be expressed simply as a set of my prior conditional credences in X given the evidence E_i . If I now rationally comply with an EV-maximizing decision rule in light of my new representor, I will be vulnerable to a diachronic Dutch book in this case (by the reasoning underlying *Proposition 3*). But it's unclear whether this kind of exploitability is a mark of irrationality. In this case, since I lose evidence and comply with other constraints of epistemic and instrumental rationality, I am unable to comply with a decision rule that makes me invulnerable to exploitation. But the *Principle of Exploitability* requires an agent to comply with a decision rule that makes them invulnerable to predictable exploitation *if* they are able to comply with that rule. So, my predictable exploitability in this case won't necessarily indicate irrationality.

Finally, consider a case where (c) is true. This will be a case where partitionality fails: my posterior evidence after observing the coin flips is either not factive (i.e., entails falsehoods), or lacks perfect introspective access itself. In such a case, if I update according to a conditionalizing rule and act according to an EV-maximizing decision rule, I may be predictably vulnerable to a diachronic Dutch book.³⁶ But, once again, it's unclear whether this kind of exploitability indicates irrationality. Suppose that updating according to *Generalized Conditionalization* and complying with an EV-maximizing rule are requirements of rationality. When partitionality fails, an epistemically rational agent who updates by conditionalization will update either on a falsehood or on misleading evidence about their own evidence. If the agent is instrumentally rational in other respects and therefore complies with an EV-maximizing rule, they will be unable to comply with any decision rule that could prevent them from being predictably exploitable. Since the agent is unable to comply with such a decision rule in virtue of being epistemically and instrumentally rational in other respects, predictable exploitability in cases of this sort won't necessarily indicate irrationality.

The result is this. According to the *Principle of Exploitability*, if an agent is able to comply with a decision rule, then it is instrumentally rational for that agent to comply with a decision rule only if it doesn't make them predictably exploitable. In cases where *Practical Stability* doesn't prevent an agent (who satisfies other requirements of epistemic and instrumental rationality) from being exploitable, the relevant exploitability is either unpredictable or due to the agent's inability to comply with a decision rule that will prevent them from being exploitable (provided that the agent is epistemically and instrumentally rational in other respects). Neither kind of exploitability—according to *Principle of Exploitability*—is indicative of irrationality.

credences to the possibilities compatible with my evidence.

³⁶For recent discussions of this point, see Gallow (2019) and Das (2022).

13 Challenge 2: Collapse

The second challenge targets a feature of *Practical Stability* that we have already considered. *Practical Stability* requires an agent with an imprecise doxastic state to pick a probability function from their representor, and to use that probability function—and conditionalized versions of that function—for making decisions within the context of an inquiry (provided that they do so). Thus, it requires an agent with an imprecise doxastic state make choices within the context of an inquiry in the same way in which an agent with a precise doxastic state would.

The problem is this. If any two inquiries within an agent's life could be treated as part of a single partitional inquiry, then any segment of their life—stretching from the present to a future moment—could be represented as a single inquiry. Then, *Practical Stability* would say that they should (if they can) just pick a probability function from their current rational representor and make decisions throughout that inquiry using that function or suitably conditionalized versions of that function. An agent who complies with this constraint would make choices exactly like a rational agent with a precise doxastic state. So, there would be no functional difference between the doxastic state of this agent and a precise doxastic state. Suppose we accept functionalism about doxastic states and, therefore, think that doxastic states are individuated in terms of the functional role that they play in relation to other mental states and behaviour. Then, we will be compelled to say that a rational agent whose doxastic state we were originally treating as imprecise in fact has the same doxastic state as a rational agent with a precise doxastic state (since it interacts with other mental states and influences behaviour in the same way). Thus, the worry is that *Practical Stability* collapses the distinction between rational agents with precise doxastic states and rational agents with imprecise ones.³⁷

This worry is based on the assumption that it is possible to treat any segment of an agent's life as a partitional inquiry. But that's implausible. Of course, there may be ideal agents who can anticipate all the possible evidence-gathering acts or events that lie ahead of them, or who never lose any evidence as they gain more evidence, or whose evidence is always factive and always has perfect introspective access to itself. We are not like that. We often cannot conceive—let alone rationally anticipate—which evidence-gathering acts or events we may perform or undergo in the future. We often lose evidence, e.g., through forgetting, in the course of acquiring new evidence. And we often lack perfect access to our own evidence. So, not every segment of our lives—from the present to some future moment—can be treated as a single partitional inquiry. In fact, we are best represented as agents who, at different stages of our lives, engage in distinct inquiries that cannot be embedded within a single partitional inquiry. That is why *Practical Stability* doesn't require us to comply with any decision rule that remains

³⁷Thanks to Brian Weatherson for suggesting this worry.

stable throughout our lives even if we are able to do so. So, it preserves—at least for non-ideal agents like us—the distinction between rational agents with precise doxastic states and rational agents with imprecise doxastic states.³⁸

14 Challenge 3: Against the Dutch Book Argument

A different, but I think more serious, challenge for *Practical Stability* targets the Dutch book argument sketched in §11. Recall the diachronic Dutch book laid out in Table 3. If I comply with a stable decision rule in *Mystery Urn 2*, I will initially accept bet a , but won't necessarily accept b after my inquiry. If I accept b on learning that the coin landed heads (or tails), I must turn it down on learning that the coin landed tails (or heads). But, on reflection, this might seem arbitrary. Irrespective of what I learn about the coin flip, b will look pretty much the same to me. I will assign an interval of credences $[0, 1]$ to B and $\sim B$, and b yields a payoff of -0.8 if B is true and of 0.2 if $\sim B$ is false. Why should my future self make different choices with regard to b in these two situations, given that its attitudes towards b are the same?

Appealing to the threat of exploitation won't help here. Before the inquiry begins, if I see the diachronic Dutch book coming, I will prefer to comply with a stable decision rule, because complying with an unstable rule will make me exploitable. So, I may decide to comply with a stable decision rule in the future. But, even if my past self made such a decision, it's not obvious why that earlier decision should give my future self any reason to use a stable decision rule. Arguably, instrumental rationality is a matter of doing what is optimal in light of one's current doxastic state and preferences. If my future self doesn't care about honouring the decisions of my past self, it's not clear why my future self is required by instrumental rationality to act according to a stable decision rule and therefore act differently in two situations where it has pretty much the same attitudes with respect to bet b .³⁹ So, even in a case where an agent faces

³⁸It's worth noting that the success of this response depends partly on the background assumption that *intrapersonal permissivism*—the view that it can be rational for the same agent at different times to adopt different doxastic attitudes towards the same proposition in response to the same body of evidence—is false. If intrapersonal permissivism were true, then an agent would be rationally permitted to switch the epistemic standards by which they weigh their evidence in the context of different inquiries. We could simply represent a rational agent whose doxastic state can be represented by a set of credence functions as a rational agent who has a precise doxastic state at any time but changes her epistemic standards across different inquiries. This would help one resurrect the collapse challenge. I have two brief responses. First, our initial argument for the *Rationality of Imprecision* in §1 presupposed that permissivism is false (see fn. 7). So, if one found that argument plausible, one shouldn't resurrect the collapse challenge by appealing to permissivism. Second, even if permissivism is true, it seems to me—given arguments offered by White (2005) and Schultheis (2018)—that intrapersonal permissivism is much less plausible than *interpersonal* permissivism, the view that it can be rational for the two agents to adopt different doxastic attitudes towards the same proposition in response to the same body of evidence. If one accepts this, one shouldn't rely on interpersonal permissivism to resurrect the collapse challenge.

³⁹A version of this worry is raised by Adam Elga (2010) in response to what he calls “the planning proposal,” the proposal according to which—in situations where an agent anticipates that she might act in a suboptimal manner unless otherwise constrained—they should make a plan not to act in those ways and stick to it in the future. This proposal is closely related to a stronger view—endorsed by *resolute choice*

a diachronic Dutch book, they may not be required by instrumental rationality to act according to a stable decision rule *throughout* their inquiry.

It's worth noticing that this argument is motivated by a restrictive conception of instrumental rationality that is incompatible with the *Principle of Exploitability*. The *Principle of Exploitability* says that, if an agent is able to comply with a decision rule, they are permitted by instrumental rationality to comply with it only if it doesn't make them predictably exploitable. On the restrictive conception of instrumental rationality, there is no such connection between instrumental irrationality and predictable exploitability. Instrumental rationality is a matter of doing what is optimal according to one's current doxastic state and preferences. So, even though acting according to an unstable decision rule makes an agent predictably vulnerable to exploitation, it can be instrumentally rational for them to comply with such a rule at later stages of their inquiry as long as doing so is optimal according to their doxastic states and preferences at those times.

It seems to me that the requirements of instrumental rationality aren't quite as narrow as this view makes them out to be. Instrumental rationality requires us not only to choose acts that are optimal given our current doxastic states and preferences, but also to cultivate practical dispositions that help us escape predictable exploitation. So, consider:

The Principle of Exploitability for Dispositions. If making choices of a certain pattern makes an agent predictably exploitable, then the agent is required by instrumental rationality to be disposed not to make choices of that pattern.

This principle arguably underlies the standard money arguments in favour of putative constraints of instrumental rationality like the requirement of having transitive preferences. For instance, the money pump argument for the transitivity of preferences starts out with the observation that choosing according to intransitive preferences can make an agent predictably exploitable. Given this, the *Principle of Exploitability for Dispositions* says that an agent is required by instrumental rationality to be disposed not to choose as an agent with intransitive preferences would. But, arguably, preferences are just dispositions to choose in certain ways. This means that instrumental rationality requires us not to have intransitive preferences. Since the *Principle of Exploitability for Dispositions* helps us vindicate independently plausible constraints of instrumental rationality on preferences, it seems quite plausible.

Turn now to the case of unstable decision rules. Since complying with unstable rules makes us predictably exploitable, the *Principle of Exploitability for Dispositions* will predict that we are required by instrumental rationality to be disposed not to comply

theorists like Gauthier (1986, 1997, 1998) and McClennen (1990, 1997)—according to which, if an agent rationally adopts a plan in light of her earlier preferences, then they are later rationally required to stick to the plan even if that involves acting against their current preferences. I don't endorse any version of this planning proposal.

with such unstable rules. But, next, consider a further assumption:

The Disposition-Choice Principle. If an agent

- is required by instrumental rationality to be disposed not to comply with a certain decision rule R within the context of an inquiry, and
- is able to comply with some other decision rule R^* ,

then they are required by instrumental rationality not to comply with R within the context of an inquiry.

This principle is motivated by a simple thought about the connection between practical dispositions and the choices that manifest them. What does it mean for an agent to be disposed not to make certain choices under certain circumstances? It just means that the agent won't make those choices under those circumstances, provided that they are able to make other choices. So, if an agent is required by instrumental rationality to be disposed not to comply with a certain decision rule within the context of an inquiry (because it makes them predictably exploitable), then they are required by instrumental rationality not to comply with it, provided that they are able to comply with some other rule. This supports the *Principle of Exploitability*. And this implies that, if an agent is able to comply with a stable decision rule in a partitional inquiry, they are required by instrumental rationality not to comply with an unstable decision rule in that inquiry.

The challenge for someone who rejects the *Principle of Exploitability* is this: they will have to reject either the *Principle of Exploitability for Dispositions* or the *Disposition-Choice Principle*. If they reject *Principle of Exploitability for Dispositions*, it will be difficult for them to vindicate the standard money pump arguments for the traditionally accepted constraints of instrumental rationality on preferences. If they give up the *Disposition-Choice Principle*, they will be breaking a natural tie between practical dispositions and the choices that manifest them.

Even if this line of reasoning doesn't ultimately succeed, the concept of stable decision rules remains theoretically fruitful. Let's reject *Practical Stability*, but accept *Caprice*, the norm that an agent is permitted by instrumental rationality to perform an act just in case it maximizes expected value by lights of some member of their current representor. Then, given a few other assumptions, *Proposition 2* will entail the following principle:

Moderately Restricted Value of Evidence. Suppose a piece of evidence is available to an agent for gathering and use at a negligible cost through a partitional inquiry, and the agent is rationally and correctly certain that they will respond to that evidence in an epistemically and instrumentally rational manner *by complying with a practically stable decision rule*. Then, it is instrumentally rational for that agent to gather that evidence and use it for making decisions.

Proposition 2 says that, if an agent is rationally certain that they will act according to a stable decision rule in a partitional inquiry, then Good’s inequality holds. This, in turn, will guarantee that, by lights of some member of their prior representor, it will be optimal for the agent to gather and use the available evidence (provided that it’s cost-free). So, by *Caprice*, it will follow that it’s instrumentally rational for the agent to do so. Thus, the discussion of stable decision rules sheds light on the conditions under which *Value of Evidence* holds for agents with imprecise doxastic states.

15 Conclusion

It’s time to take stock. In this paper, I have explored a tension between *Value of Evidence* and *Rationality of Imprecision*. As promised, I have shown three things.

First, I have shown that the tension emerges because *Rationality of Imprecision*—together with *Generalized Conditionalization*—gives rise to unreflective inquiries. In such unreflective inquiries, synchronic norms of instrumental rationality for agents with imprecise doxastic states lead to failures of *Value of Evidence*.

Second, I have shown that we can resolve the tension simply accepting a diachronic norm of rationality called *Practical Stability*. Roughly, this says that, in the context of a partitional inquiry, an agent is permitted to act according to a decision rule just in case it is stable, i.e., recommends acts which maximize expected value relative to probability functions that diachronically cohere in the context of that inquiry.

Third, I have shown that *Practical Stability* can be partially motivated by means of a Dutch book argument, and can be defended against potential objections. And, even if it is ultimately indefensible, our discussion of this norm helps us more clearly understand the conditions under which *Value of Evidence* holds for agents with imprecise doxastic states.⁴⁰

Appendix A: Good’s Theorem

Good’s Theorem. For any partitional inquiry $\langle W, E, \Gamma \rangle$, let Γ contain a single probability function p . If the agent is rationally certain that they will update according to a conditionalizing rule U and comply with an EV-maximizing decision rule R for any decision problem, then, for any decision problem $D = \langle W, A, v \rangle$,

$$\sum_{w \in W} p(w)v(R(D, \Gamma), w) \leq \sum_{w \in W} p(w)v(R(D, U(w)), w).$$

And, if, for any $w \in W$, $R(D, \Gamma)$ doesn’t maximize expected value relative to any $p^* \in U(w)$, then

⁴⁰I wish to thank Mike Caie, Cian Dorr, Verónica Gómez-Sánchez, Harvey Lederman, Ram Neta, Laurie Paul, Jeff Russell, Ted Sider, and others present at the 2022 Ranch Metaphysics Workshop, for their comments on this paper. I am especially grateful to Joe Horton for numerous discussions of these themes, to Brian Weatherson for written comments, and to two anonymous referees for this journal for their helpful suggestions for improving the paper.

$$\sum_{w \in W} p(w)v(R(D, \Gamma), w) < \sum_{w \in W} p(w)v(R(D, U(w)), w).$$

Proof. Let the set of all the strongest pieces of evidence that the agent could get in the inquiry $\langle W, E, \Gamma \rangle$ be $\mathcal{E} = \{E_1, \dots, E_k\}$. We know that R is an EV-maximizing rule, and Γ is a singleton set of probability functions containing some probability function p . We can write the left hand side of the inequality as:

$$\begin{aligned} \sum_{w \in W} p(w)v(R(D, \Gamma), w) &= \sum_{i=1}^k \sum_{w \in E_i} p(E_i)p(w|E_i)v(R(D, \Gamma), w) \\ &= \sum_{i=1}^k \sum_{w \in E_i} p(w)v(R(D, \Gamma), w) \\ &= \max_{a \in A} \sum_{i=1}^k \sum_{w \in E_i} p(w)v(a, w) \end{aligned}$$

And we write the right hand side of the inequality as:

$$\begin{aligned} \sum_{w \in W} p(w)v(R(D, U(w)), w) &= \sum_{i=1}^k \sum_{w \in E_i} p(E_i)p(w|E_i)v(R(D, U(w)), w) \\ &= \sum_{i=1}^k \sum_{w \in E_i} p(w)v(R(D, U(w)), w) \\ &= \sum_{i=1}^k \max_{a \in A} \sum_{w \in E_i} p(w)v(a, w) \end{aligned}$$

So, the whole inequality becomes:

$$\max_{a \in A} \sum_{i=1}^k \sum_{w \in E_i} p(w)v(a, w) \leq \sum_{i=1}^k \max_{a \in A} \sum_{w \in E_i} p(w)v(a, w).$$

To prove this inequality, we only need the following lemma.

Lemma. $\sum_k \max_j f(j, k) \geq \max_j \sum_k f(j, k)$, with strict inequality unless $\max_j f(j, k)$ is satisfied by one value of j for any k .

Let a value of j that maximizes $\sum_k f(j, k)$ be j_{max} . Now, for any k , $\max_j f(j, k) \geq f(j_{max}, k)$, irrespective of how we define j_{max} . The inequality is strict unless j_{max} is a value that maximizes not only $\sum_k f(j, k)$, but also maximizes $f(j, k)$ for all k . If this were to happen, the value of $\max_j f(j, k)$ would be equal to $\max_j \sum_k f(j, k)$, for any k , and therefore would be the same, for any k . Therefore, the *Lemma* is proved, and so is *Good's Theorem*. \square .

Appendix B: Proofs

Proposition 1. For any inquiry $\langle W, E, \Gamma \rangle$, suppose the relevant agent is rationally certain that they will update according to an unreflective rule U . Then, the following claims are true.

- *Claim 1.1.* For any decision rule R that satisfies Γ -Maximin, Good's inequality fails for R and some decision problem D .
- *Claim 1.2.* There exists a decision rule R that satisfies E -Admissibility and Maximality such that Good's inequality fails for R and some decision problem D .

Proof. If U is unreflective, then $\inf \Gamma(X) > \inf U(w)(X)$ for any w in W . Consider the set $\text{Min}(X) = \{\inf U(w)(X) : w \in W\}$, which is the set of all the minimum credences that X gets assigned by the agent's posterior representor in all the worlds. We pick a number r that lies between $\inf \Gamma(X)$ and the greatest element in $\text{Min}(X)$ (exclusive). Using r , we can create a decision problem $D = \langle W, A, v \rangle$ such that $A = \{\text{Accept}, \text{Reject}\}$ and the values of v are given by Table 4. Let R be any action satisfies Γ -Maximin.

	Y	$\sim Y$
<i>Accept</i>	$\frac{1-r}{r}$	-1
<i>Reject</i>	0	0

Table4: A Payoff Table for the Proof of *Proposition 1*

For any $p \in \Gamma$, the expected value of *Accept* is:

$$\begin{aligned}
 \text{Exp}_p(\text{Accept}) &= p(X) \cdot \left(\frac{1-r}{r}\right) + p(\sim X) \cdot (-1) \\
 &= \frac{p(X) - r \cdot p(X) - r \cdot p(\sim X)}{r} \\
 &= \frac{p(X) - r}{r}
 \end{aligned}$$

Since $p(X)$ is greater than r , the the expected value of *Accept* is positive. Since R satisfies Γ -Maximin, $R(D, \Gamma) = \text{Accept}$.

In any world w , let there be a probability function p^* in $U(w)$, such that $p^*(X) = \inf U(w)(X)$. By the same calculation as above, the expected value of *Accept* according to p^* is $\frac{p^*(X) - r}{r}$. Since $p^*(X)$ is smaller than r , the expected value of *Accept* is negative. For any other probability function in $U(w)$, the expected value of *Accept* is higher. Since R satisfies Γ -Maximin, $R(D, U(w)) = \text{Reject}$.

Finally, for any $p \in \Gamma$,

$$\begin{aligned}
\sum_{w \in W} p(w)v(R(D, \Gamma), w) &= \sum_{w \in W} p(w)v(\text{Accept}, w) \\
&= \frac{p(X) - r}{r} \\
&> 0 \\
&= \sum_{w \in W} p(w)v(\text{Reject}, w) \\
&= \sum_{w \in W} p(w)v(R(D, U(w)), w).
\end{aligned}$$

So, Good's inequality fails. This completes our proof of *Claim 1.1*. But since R is also an EV-maximizing rule, this also shows that *Claim 1.2* is true. \square

Corollary 1. For any partitional inquiry $\langle W, E, \Gamma \rangle$, suppose the relevant agent is rationally certain that they will update according to an unreflective rule U and act according to a decision rule R . Then, if R satisfies *Levi's Rule*, then Good's inequality fails for R and some decision problem D .

Proof. To see why this is true, let R be a decision rule that satisfies Levi's rule. If U is unreflective, then there is some proposition X such that, for any w , $\inf U(w)(X) < \inf \Gamma(X)$. Once again, we can construct the same decision problem whose payoff table is given in Table 4. We know that, relative to each probability function in Γ , there is only one option that maximizes expected value: *Accept*. By contrast, for any w , while *Reject* is certainly E-admissible relative to some probability function in $U(w)$, there may be other probability functions in $U(w)$ relative to which *Accept* maximizes expected value. Notice that the set of expected values for *Reject* is $\{0\}$. But, if p^* is a probability function $U(w)$ be such that $p^*(X) = \inf U(w)(X)$, then $\frac{p^*(X) - r}{r}$ will be in the set of expected values for *Accept*. Since this is negative, *Reject* will maximize minimum expected value. So, R will recommend *Reject*. Thus, Good's inequality will fail for R and this decision problem. \square

Corollary 2. For any partitional inquiry $\langle W, E, \Gamma \rangle$, suppose the relevant agent is rationally certain that they will update according to an unreflective rule U . Then, there is a decision rule R that satisfies the *Hurwicz Criterion* for some α between 0.5 and 1 (exclusive), such that Good's inequality fails for R and some decision problem D .

Proof. If U is unreflective, then, for any w , $\inf U(w)(X) < \inf \Gamma(X)$. We can find a positive real α between 0.5 and 1 (exclusive), such that, for any w , $\alpha \cdot (\inf U(w)(X)) + (1 - \alpha) \cdot (\sup U(w)(X)) < \inf \Gamma(X)$. Let R be a decision rule that recommends acts that maximize Hurwicz expected value using α . Letting r be a real number between

$\alpha \cdot (\inf U(w)(X)) + (1 - \alpha) \cdot (\sup U(w)(X))$ and $\inf \Gamma(X)$ (exclusive), we can construct the same decision problem whose payoff table is given in Table 4. We know that, relative to each probability function in Γ , there is only one option that maximizes expected value: *Accept*. So, this option will also maximize Hurwicz expected value. By contrast, for any w , let p_1, p_2 be two probability functions in $U(w)$ such that $p_1(X) = \inf U(w)(X)$ and $p_2(X) = \sup U(w)(X)$. So, the Hurwicz expected value of *Accept* will be:

$$\begin{aligned}
& HExp_{U(w), \alpha}(Accept) \\
&= \alpha \cdot (\inf \{Exp_p(Accept) : p \in U(w)\}) + (1 - \alpha) \cdot (\sup \{Exp_p(Accept) : p \in U(w)\}) \\
&= \alpha \cdot (p_1(X) \cdot (\frac{1-r}{r}) + p_1(\sim X) \cdot (-1)) + (1 - \alpha) \cdot (p_2(X) \cdot (\frac{1-r}{r}) + p_2(\sim X) \cdot (-1)) \\
&= \alpha \cdot (\frac{p_1(X) - r}{r}) + (1 - \alpha) \cdot (\frac{p_2(X) - r}{r}) \\
&= \frac{\alpha \cdot p_1(X) - r + (1 - \alpha) \cdot p_2(X) - (1 - \alpha) \cdot r}{r} \\
&= \frac{\alpha \cdot p_1(X) + (1 - \alpha) \cdot p_2(X) - r}{r} \\
&< 0.
\end{aligned}$$

In contrast, $HExp_{U(w), \alpha}(Reject) = 0$. For any w , R will recommend *Reject* relative to the agent's posterior doxastic state. Thus, Good's inequality will fail for R and the relevant decision problem. \square

Proposition 2. For any partitional inquiry $\langle W, E, \Gamma \rangle$, any p in Γ , any p -stable decision rule R and any decision problem D , Good's inequality will hold relative to p .

Proof. *Proposition 2* has the same proof as *Good's Theorem*. \square

Proposition 3. For any partitional inquiry $\langle W, E, \Gamma \rangle$, suppose the agent is rationally certain that they will update according to a conditionalizing rule U . Then, if R is an EV-maximizing rule, it is not subject to a diachronic Dutch book iff it is p -stable for some p in Γ .

Proof. Let the set of strongest possible pieces of evidence that the agent could get in this inquiry be $\mathcal{E} = \{E_1, \dots, E_k\}$.

First, we shall show that, if R is EV-maximizing but not p -stable for any p in Γ , then it is subject to a diachronic Dutch book. R can be EV-maximizing but not p -stable for any p in Γ only if: (i) there is a probability function $p_1 \in \Gamma$ such that, relative to the agent's prior doxastic state, R recommends acts that maximize expected value relative to p_1 for any decision problem D , but (ii) there is some $p_2 \in \Gamma$, such that, for some evidence proposition E_i and some proposition X , $p_2(X|E_i) < p_1(X|E_i)$, and, in any E_i -world, R recommends acts that maximize expected value relative to $p_2(\cdot|E_i)$ for any decision problem. We can now show that R is subject to a diachronic Dutch book.

Let δ and ϵ be two real numbers such that:

$$(1) 0 < \delta < \frac{p_1(X|E_i) - p_2(X|E_i)}{2p_1(E_i)}.$$

$$(2) \epsilon = p_1(X \cap E_i) + p_1(E_i)(p_1(X|E_i) - p_2(X|E_i)) + p_1(\sim E_i)p_1(X|E_i) - \delta.$$

Now, consider three bets.

	$X \cap E_i$	$\sim X \cap E_i$	$\sim E_i$
b	$1 + p_1(X E_i) - p_2(X E_i) - \epsilon$	$p_1(X E_i) - p_2(X E_i) - \epsilon$	$p_1(X E_i) - \epsilon$
b_i	$-1 + p_2(X E_i) + \delta$	$p_2(X E_i) + \delta$	0
b_j	0	0	δ

Table5: A Diachronic Dutch Book for the Proof of *Proposition 3*

Let b be offered to the agent before their inquiry. Let b_i be offered to them after their inquiry when they learn E_i . And, finally, let b_j be offered to them after their inquiry when they learn any other E_j . Notice three things.

- First, the expected value of b relative to p_1 is:

$$\begin{aligned} \text{Exp}_{p_1}(b) &= p_1(X \cap E_i)(1 + p_1(X|E_i) - p_2(X|E_i) - \epsilon) \\ &\quad + p_1(\sim X \cap E_i)(p_1(X|E_i) - p_2(X|E_i) - \epsilon) + p_1(\sim E_i)(p_1(X|E_i) - \epsilon) \\ &= p_1(X \cap E_i) + p_1(E_i)(p_1(X|E_i) - p_2(X|E_i)) + p_1(\sim E_i)(p_1(X|E_i) - \epsilon) \\ &= \epsilon + \delta - \epsilon \\ &= \delta. \end{aligned}$$

So, since the expected value of b is positive, the agent will accept b .

- Second, the expected value of b_i relative to $p_2(\cdot|E_i)$ is:

$$\begin{aligned} \text{Exp}_{p_2(\cdot|E_i)}(b_i) &= p_2(X|E_i)(-1 + p_2(X|E_i)) + \delta + p_2(\sim X|E_i)(p_2(X|E_i) + \delta) \\ &= -p_2(X|E_i)(1 - p_2(X|E_i)) + p_2(X|E_i)(1 - p_2(\sim X|E_i)) + \delta \\ &= \delta \end{aligned}$$

So, the agent should accept the bet.

- Finally, the expected value of b_j in any $\sim E_i$ -world (for any j other than i) is δ . So, the agent should accept the bet.

As a result, the agent will be subject to a net loss of utility in every case.

- In $\sim E_i$ -worlds, the net utility that the agent can get is:

$$\begin{aligned}
& p_1(X|E_i) - \epsilon + \delta \\
&= p_1(X|E_i) - (p_1(X \cap E_i) + p_1(E_i)(p_1(X|E_i) - p_2(X|E_i)) + p_1(\sim E_i)p_1(X|E_i) - \delta) + \delta \\
&= 2\delta + p_1(X|E_i) - p_1(X \cap E_i) - p_1(E_i)p_1(X|E_i) + p_1(E_i)p_2(X|E_i) - p_1(\sim E_i)p_1(X|E_i) \\
&= 2\delta + p_1(X|E_i)(1 - p_1(\sim E_i)) - p_1(X \cap E_i) - p_1(E_i)p_1(X|E_i) + p_1(E_i)p_2(X|E_i) \\
&= 2\delta + p_1(X|E_i)p_1(E_i) - p_1(X \cap E_i) - p_1(E_i)p_1(X|E_i) + p_1(E_i)p_2(X|E_i) \\
&= 2\delta - p_1(X|E_i)p_1(E_i) + p_1(E_i)p_2(X|E_i) \\
&< 0.
\end{aligned}$$

- Similarly, in $X \cap E_i$ -worlds, the net utility that the agent can get is:

$$\begin{aligned}
& (1 + p_1(X|E_i) - p_2(X|E_i) - \epsilon) + (-1 + p_2(X|E_i) + \delta) \\
&= p_1(X|E_i) - \epsilon + \delta \\
&< 0.
\end{aligned}$$

- Finally, in $\sim X \cap E_i$ -worlds, the net utility that the agent can get is:

$$\begin{aligned}
& p_1(X|E_i) - p_2(X|E_i) - \epsilon + p_2(X|E_i) + \delta \\
&= p_1(X|E_i) - \epsilon + \delta \\
&< 0.
\end{aligned}$$

So, R is subject to a diachronic Dutch book. Therefore, if R is unstable but EV-maximizing, then it is subject to a diachronic Dutch book.

Next, we will show that, if R is p -stable (for some p in Γ), then it's not subject to a diachronic Dutch book. Suppose R is p -stable (for some p in Γ), So, p is the probability function such that, for any decision problem D , $R(D, \Gamma) = a$ just in case a maximizes expected value relative to p , and, for any world w and any decision problem D , $R(D, U(w)) = a$ just in case a maximizes expected value relative to $p(\cdot|E(w))$. Suppose, for *reductio*, R is a subject to a diachronic Dutch book. So there is a series of bets b, b_1, \dots, b_k such that (i) b is offered to the agent before their inquiry and each b_i is offered just in case the agent receives E_i as their posterior evidence, and (iii) relative to her prior and posterior doxastic state, the agent will accept all the bets, and will undergo a net loss of utility. Let D be the decision problem that involves a choice between accepting b and rejecting b , and, for each b_i , let D_i be the decision problem that involves a choice between accepting b_i and rejecting b_i . Then, $R(D, \Gamma) = b$ and, for any E_i -world, $R(D_i, U(w)) = b_i$.

We know that, for any E_i -world, accepting both b and b_i results in a net loss of

utility. So, we know that:

$$\sum_{i=1}^k \sum_{w \in E_i} p(w)(v(b, w) + v(b_i, w)) < 0.$$

Given that R is an EV-maximizing rule, $Exp_p(b) \geq 0$ and $Exp_{p(\cdot|E_i)}(b_i) \geq 0$ for i between 1 and k (inclusive). But notice that the sum of $Exp_p(b)$ and the expectation of the expected values of the b_i 's is:

$$\begin{aligned} & \sum_{w \in W} p(w)v(b, w) + \sum_{i=1}^k p(E_i) \sum_{w \in W} p(w|E_i)v(b_i, w) \\ &= \sum_{w \in W} p(w)v(b, w) + \sum_{i=1}^k p(E_i) \sum_{w \in E_i} p(w|E_i)v(b_i, w) \\ &= \sum_{w \in W} p(w)v(b, w) + \sum_{i=1}^k \sum_{w \in E_i} p(E_i)p(w|E_i)v(b_i, w) \\ &= \sum_{i=1}^k \sum_{w \in E_i} p(w)v(b, w) + \sum_{i=1}^k \sum_{w \in E_i} p(w)v(b_i, w) \\ &= \sum_{i=1}^k \sum_{w \in E_i} p(w)(v(b, w) + v(b_i, w)). \end{aligned}$$

So, this cannot be less than 0, since $Exp_p(b) \geq 0$ and $Exp_{p(\cdot|E_i)}(b_i) \geq 0$ for each i between 1 and k (inclusive). So, we have arrived at a contradiction. Therefore, R cannot both be a p -stable decision rule and be subject to a diachronic Dutch book. \square

References

- Ahmed, A. (2014). *Evidence, Decision and Causality*. Cambridge: Cambridge University Press.
- Ahmed, A. and Salow, B. (forthcoming). Don't look now. *The British Journal for the Philosophy of Science*.
- Berger, J.O. (1985). *Statistical decision theory and Bayesian analysis*. Springer-Verlag, New York.
- Blackwell, D. (1951). Comparison of experiments. In *Proceedings, Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. Berkeley: University of California Press.
- Bradley, S. and Steele, K. (2016). Can free evidence be bad? value of information for the imprecise probabilist. *Philosophy of Science*, 83(1):1–28.

- Bradley, S. and Steele, K.S. (2014). Should subjective probabilities be sharp? *Episteme*, 11(3):277–289.
- Buchak, L. (2010). Instrumental rationality, epistemic rationality, and evidence-gathering. *Philosophical Perspectives*, 24(1):85–120.
- Campbell-Moore, C. and Salow, B. (2020). Avoiding risk and avoiding evidence. *Australasian Journal of Philosophy*, 98(3):495–515.
- Das, N. (2020). The value of biased information. *British Journal for the Philosophy of Science*, page axaa003.
- Das, N. (2022). Externalism and exploitability. *Philosophy and Phenomenological Research*, 104(1):101–128.
- Davidson, D., McKinsey, J. C.C., and Suppes, P. (1955). Outlines of a formal theory of value, i. *Philosophy of Science*, 22(2):140–160.
- DeFinetti, B. (1937/1992). Foresight: Its logical laws, its subjective sources. In H.E.Kyburg, . H. E. K.S., editor, *Studies in Subjective Probability*, pages 134–174. Huntington, NY: Robert E. Kreiger Publishing Co.
- Dorst, K. (forthcoming). Evidence: A guide for the uncertain. *Philosophy and Phenomenological Research*.
- Douven, I. (2009). Uniqueness revisited. *American Philosophical Quarterly*, 46(4):347–361.
- Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers' Imprint*, 10.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, pages 643–669.
- Eva, B. (2019). Principles of indifference. *Journal of Philosophy*, 116(7):390–411.
- Gallow, J.D. (2019). Diachronic dutch books and evidential import. *Philosophy and Phenomenological Research*, 99(1):49–80.
- Gauthier, D. (1997). Resolute choice and rational deliberation: A critique and a defense. *Noûs*, 31(1):1–25.
- Gauthier, D. (1998). Rethinking the toxin puzzle. In Coleman, J.L., Morris, C.W., and Kavka, G.S., editors, *Rational Commitment and Social Justice: Essays for Gregory Kavka*, pages 47–58. Cambridge University Press.
- Gauthier, D.P. (1986). *Morals by Agreement*. Oxford: Oxford University Press.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.

- Good, I.J. (1967). On the principle of total evidence. *British Journal for the Philosophy of Science*, 17(4):319–321.
- Good, I.J. (1974). A little learning can be dangerous. *British Journal for the Philosophy of Science*, 25(4):340–342.
- Greaves, H. and Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459):607–632.
- Greco, D. and Hedden, B. (2016). Uniqueness and metaepistemology. *Journal of Philosophy*, 113(8):365–395.
- Grünwald, P.D. and Dawid, A.P. (2004). Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *The Annals of Statistics*, 32(4):1367 – 1433.
- Hart, C. and Titelbaum, M.G. (2015). Intuitive dilation? *Thought: A Journal of Philosophy*, 4(4):252–262.
- Hawthorne, J., Landes, J., Wallmann, C., and Williamson, J. (2017). The principal principle implies the principle of indifference. *British Journal for the Philosophy of Science*, 68(1).
- Horowitz, S. (2014). Immoderately rational. *Philosophical Studies*, 167(1):41–56.
- Horowitz, S. (2019). The truth problem for permissivism. *Journal of Philosophy*, 116(5):237–262.
- Hurwicz, L. (1951). The generalized bayes minimax principle: a criterion for decision making under uncertainty. *Cowles Comm. Discuss. Paper Stat*, 335:1950.
- Jaynes, E.T. (1957). Information theory and statistical mechanics. *Physical Review*, 106:620–630.
- Joyce, J.M. (2005). How probabilities reflect evidence. *Philosophical Perspectives*, 19(1):153–179.
- Joyce, J.M. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives*, 24:281–323.
- Kadane, J.B., Schervish, M., and Seidenfeld, T. (2008). Is ignorance bliss? *Journal of Philosophy*, 105(1):5–36.
- Kaplan, M. (1983). Decision theory as philosophy. *Philosophy of Science*, 50(4):549–577.
- Keynes, J.M. (1921). *A Treatise on Probability*. London: Macmillan and Co.
- Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy*, 71(13):391–418.

- Levi, I. (1980). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. Cambridge, MA: MIT Press.
- Levi, I. (1986). *Hard Choices: Decision Making Under Unresolved Conflict*. Cambridge University Press.
- Lewis, D.K. (1980). A subjectivist's guide to objective chance. In Jeffrey, R.C., editor, *Studies in Inductive Logic and Probability, Volume II*, pages 263–293. Berkeley: University of California Press.
- McClennen, E.F. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge: Cambridge University Press.
- Mcclennen, E.F. (1997). Pragmatic rationality and rules. *Philosophy and Public Affairs*, 26(3):210–258.
- Moss, S. (2015). Time-slice epistemology and action under indeterminacy. *Oxford Studies in Epistemology*, 5:172–94.
- Moss, S. (2020). Global constraints on imprecise credences: Solving reflection violations, belief inertia, and other puzzles. *Philosophy and Phenomenological Research*.
- Pedersen, A.P. and Wheeler, G. (2014). Demystifying dilation. *Erkenntnis*, 79(6):1305–1342.
- Pedersen, A.P. and Wheeler, G. (2015). Dilation, disintegrations, and delayed decisions. In Augustin, T., Dora, S., Miranda, E., and Quaeghebeur, E., editors, *Proceedings of the 9th International Symposium on Imprecise Probability: Theories and Applications (ISIPTA 2015)*, pages 227–236. Aracne Editrice.
- Peirce, C.S. (1967). Note on the theory of the economy of research. *Operations Research*, 15(4):643–648.
- Rabinowicz, W. (2009). Letters from long ago: On causal decision theory and centered chances. In Lars-Göran, J., editor, *Logic, Ethics, and All That Jazz-Essays in Honour of Jordan Howard Sobel*. Uppsala: Uppsala Philosophical Studies.
- Ramsey, F. (1926). Truth and probability. In Eagle, A., editor, *Philosophy of Probability: Contemporary Readings*, pages 52–94. Routledge.
- Ramsey, F.P. (1990). Weight or the value of knowledge. *The British Journal for the Philosophy of Science*, 41(1):1–4.
- Rinard, S. (2013). Against radical credal imprecision. *Thought: A Journal of Philosophy*, 2(1):157–165.

- Salow, B. (2018). The externalist's guide to fishing for compliments. *Mind*, 127(507):691–728.
- Schick, F. (1986). Dutch bookies and money pumps. *Journal of Philosophy*, 83(2):112–119.
- Schoenfield, M. (2014). Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*, 48(2):193–218.
- Schoenfield, M. (2017). Conditionalization does not maximize expected accuracy. *Mind*, 126(504):1155–1187.
- Schoenfield, M. (2019). Permissivism and the value of rationality: A challenge to the uniqueness thesis. *Philosophy and Phenomenological Research*, 99(2):286–297.
- Schultheis, G. (2018). Living on the edge: Against epistemic permissivism. *Mind*, 127(507):863–879.
- Seidenfeld, T. (2004). A contrast between two decision rules for use with (convex) sets of probabilities: Γ -maximin versus e-admissibility. *Synthese*, 140(1-2):69–88.
- Skyrms, B. (1990). The value of knowledge. *Minnesota Studies in the Philosophy of Science*, 14:245–266.
- Sturgeon, S. (2008). Reason and the grain of belief. *Noûs*, 42(1):139–165.
- Titelbaum, M.G. (2010). Not enough there there: Evidence, reasons, and language independence. *Philosophical Perspectives*, 24:477–528.
- Titelbaum, M.G. and Kopec, M. (ms). Plausible permissivism.
- Vallinder, A. (2018). Imprecise bayesianism and global belief inertia. *British Journal for the Philosophy of Science*, 69(4):1205–1230.
- van Fraassen, B.C. (1984). Belief and the will. *Journal of Philosophy*, 81(5):235–256.
- van Fraassen, B.C. (1989). *Laws and Symmetry*. Oxford: Oxford University Press.
- van Fraassen, B.C. (1990). Figures in a probability landscape. In Dunn, J. and Gupta, A., editors, *Truth or Consequences*, pages 345–356. Kluwer Academic Publishers.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman & Hall.
- Weatherson, B. (1998). Decision making with imprecise probabilities.
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19(1):445–459.
- White, R. (2009). Evidential symmetry and mushy credence. In Gendler, T.S. and Hawthorne, J., editors, *Oxford Studies in Epistemology*, pages 161–186. Oxford University Press.

Williams, J. R.G. (2014). Decision-making under indeterminacy. *Philosophers' Imprint*, 14.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.