

A forward looking decision rule for imprecise credences

Rohan Sud

Published online: 17 November 2013
© Springer Science+Business Media Dordrecht 2013

Abstract Adam Elga (*Philosophers' Imprint*, 10(5), 1–11, 2010) presents a diachronic puzzle to supporters of imprecise credences and argues that no acceptable decision rule for imprecise credences can deliver the intuitively correct result. Elga concludes that agents should not hold imprecise credences. In this paper, I argue for a two-part thesis. First, I show that Elga's argument is incomplete: there is an acceptable decision rule that delivers the intuitive result. Next, I repair the argument by offering a more elaborate diachronic puzzle that is more difficult for imprecise Bayesians to avoid.

Keywords Formal epistemology · Decision theory · Imprecise credences

Elga (2010) presents a variant of a Dutch Strategy to supporters of imprecise credences, and argues that any account of decision-making that succumbs to this book of bets is unacceptable. Elga concludes that this normative requirement, along with two additional constraints, rules out all imprecise decision theories—norms that select permissible actions for agents reasoning with imprecise credences. When we hold imprecise credences, all roads lead to irrationality, so best not to hold such credences. Or so Elga concludes.

In broad outline, Elga's argument has the following form:

1. If there is no decision rule for agents with imprecise credences that meets three constraints, then our credences should be precise.
2. There is no decision rule for agents with imprecise credences that meets these three constraints.
3. Therefore: our credences should be precise.

R. Sud (✉)
University of Michigan, Ann Arbor, MI, USA
e-mail: rsud@umich.edu

Most responses in the literature have rejected premise 1. These replies, given on behalf of the imprecise Bayesian, contest the rationality of one of Elga's three constraints (Moss ms.; Hedden 2013; McClennen 1990). In this paper, however, I will be exploring a more direct defense; I will reject premise 2. I will argue that *even if* we accept all of Elga's purported rational constraints, there remains a plausible decision rule that is up to the challenge. After the original argument is shown to be incomplete, however, the second part of the paper tests the bounds of the proposed rule. I will ultimately conclude that the proposed decision rule applies in a limited range of cases; Elga's argument against imprecise Bayesians can thus be repaired.

Here is the plan for the paper: Sect. 1 presents some background to the debate between the precise and imprecise Bayesian. Sect. 2 motivates the challenge presented in Elga (2010). Sect. 3 points out a flaw in Elga's reasoning. Formalizing this observation, I go on to propose a decision rule, *Forward Looking*, and show how it escapes Elga's criticism. In Sect. 4, I argue that *Forward Looking* is immune to simple diachronic puzzles, like Elga's original book of bets. However, I go on to propose a more elaborate diachronic puzzle that is more difficult to avoid. Sect. 5 concludes by reflecting on some broader lessons from our exploration.

1 Imprecise Bayesianism

Elga (2010) opens with the following hypothetical:

A stranger approaches you in the street and starts to pull out objects from a bag. The first three objects he pulls out are a regular-sized tube of toothpaste, a live jellyfish, and a travel-sized tube of toothpaste. To what degree should you believe that the next object he pulls out will be another tube of toothpaste?

It seems plausible that, in this case, it is not a requirement of rationality that an agent adopt a credence 'exact to the millionth decimal' in the proposition q that the next object out of the bag will be another tube of toothpaste. Instead, one might contend, an agent ought to hold credences that admit of a certain amount of imprecision. Traditionally, the credal states of agents have been represented with a single 'credence function'—a probability measure that maps propositions to a real number. Instead, we follow Joyce (2011) in representing credal states as sets of credence functions: $c^r(\varphi) = \{c_1(\varphi), c_2(\varphi), \dots\}$; we also call the credal set, c^r , the agent's *representor*. When an agent's representor, as evaluated for some proposition, results in a range of values, we say the agent has an imprecise credence with respect to that proposition. We will assume that when evaluated for any proposition, the representor yields an interval. Imprecise Bayesians defend the thesis that:

Imprecise Bayesianism There are some evidential situations in which, for some proposition φ , the representor of a perfectly rational agent may not be point-valued when evaluated for φ

In other words, they claim that the representor of a perfectly rational agent may not consist of a singleton set. Precise Bayesians reject the above thesis.

2 Elga's problem

Let's suppose that the thesis of imprecise Bayesianism has a *prima facie* appeal; at first blush, there are some evidential situations in which a rational doxastic attitude to hold is not precise. Perhaps such an appeal is generated by the evidentialist slogan that you ought to believe in accordance with your evidence: if our evidence is imprecise—as is plausibly the case when a stranger begins erratically pulling out jellyfish and toothpaste from a bag—then our credences ought to be imprecise as well. Elga, however, poses a problem for accounts that allow for such credal states. Consider an agent, Annie, whose credence in q is, say, $c'(q) = [0.1, 0.8]$. For any number in the interval $[0.1, 0.8]$, there is some credence function in Annie's representor that, when evaluated for q , yields that number. Intuitively, Annie's credence in q is 'spread out' between 0.1 and 0.8. When presented with bets on q , how should Annie evaluate such bets?

This is the challenge that Elga sets for defenders of imprecise Bayesianism: to devise a rule that imprecise agents can use to guide their decision-making. Of course, not any old decision rule will do. Elga partitions the space of possible decision rules into three classes and argues for plausible constraints that eliminate each class as candidate theories of rational choice.

2.1 Strict rules

Elga immediately sets aside a class of decision rules he calls "Strict Rules". According to such accounts, an agent ought to act just as if she has a particular precise credence function in her representor. Put more carefully, the following is a schema for such rules:

Decision Schema: Strict Specify a function, $f(\cdot)$, that maps credal sets to some particular credence function in the representor: $f(c') = c \in c'$. An act is permissible just in case it maximizes expected utility according to that credence function. (Function $f(\cdot)$ must select the same credence function in the agent's representor for all decision problems she may face.)¹

Such accounts require an imprecise agent to mimic a rational precise agent; in this way, strict rules import the standard practical consistency results of traditional expected-utility maximization. The problem with such rules, however, is that they mimic a precise agent *too* closely. Dorr (2010) puts the point succinctly:

...I think that if [Strict] is all we have to say about the decision theory [for imprecise credences], we lack an acceptable account of what it is to be in a given unsharp credential state—we cannot explain what would constitute the difference between someone in a sharp credential state given by a certain conditional probability function, and someone in an unsharp credential state

¹ Elga (2010) does not explicitly characterize Strict Rules using the schema above. Instead, he gestures towards this characterization with the *midpoint rule*. As Moss (ms., Appendix) argues, however, the midpoint rule endorses incoherent synchronic betting behavior.

containing that probability function, who had chosen it as the guide to their actions.

Elga and Dorr apparently endorse the following constraint on decision rules:

Requirement: Imprecision Given a utility function U and non-singleton representor c^r , no acceptable decision rule will permit exactly the same actions for an agent represented by U and c^r and an agent represented by U and a singleton representor, in all decision problems.

Strict decision rules are plainly in violation of *Imprecision*.

2.2 Permissive rules

If strict rules are unacceptable, one natural alternative for imprecise Bayesians is to adopt a permissive rule. According to such decision rules, there are some imprecise credal states for which the rule neither requires the agent to accept nor reject a bet.² As Elga points out, permissive rules are “very much in the spirit” of imprecise Bayesianism: “if your evidence is so unspecific as to demand a widely spread-out probability function, it is natural that the requirements of rationality be correspondingly spread out.” (p. 5). Most defenders of imprecise credences have favored permissive rules of some kind. One popular example of this type of decision rule is the following:

Decision Rule: Permissive An action is permissible just in case according to at least one probability function c in your representor c^r , the expected-utility of that action is greater than or equal to all other actions.³

Elga designs an ingenious diachronic puzzle for *Permissive*. Recall our imprecise agent Annie with belief $c^r(q) = [0.1, 0.8]$. Imagine Annie is a miser (a miser is an agent that only cares about money: her utility function increases linearly with dollars) and she is presented with the following two bets, offered sequentially at times t_1 and t_2 :

Bet 1: If q is true, you lose \$10. Otherwise you win \$15.

Bet 2: If q is true, you win \$15. Otherwise you lose \$10.

The bets are offered under the following set of conditions (“C”):

- The agent’s utility function does not change between times of the bets
- The agent does not receive evidence for the propositions being bet on between times of the bets
- The agent is rational throughout the problem

Suppose further that Annie is certain of the set-up of the problem—she is certain that she will be presented, under conditions C, with two sequential bets. According

² And these actions remain permissible even if the payoff of the bet is slightly increased.

³ This rule is a synchronic version of the *Caprice* rule defended in Weatherson (ms. a).

to *Permissive*, it is permissible for Annie to accept or reject *Bet 1*, and it is permissible for Annie to accept or reject *Bet 2*. (According to some members of Annie's representor, *Bet 1* has positive expected utility; according to others, the bet has negative expected utility. Similarly with *Bet 2*.) To see where the rule goes wrong, notice that accepting both *Bet 1* and *Bet 2* guarantees the agent \$5, whether or not q turns out to be true. Moreover, rejecting both bets leaves the agent with nothing: a payoff of \$0. Of course, any rule that permits an agent to choose a sure payoff of \$0 over a sure payoff of \$5 is unacceptable. Thus, any rule that permits an agent to choose to reject both bets is unacceptable. This is Elga's second constraint:

Requirement: Acceptance For an agent (certain of the set-up of the problem) presented with *Bet 1* and *Bet 2* (in conditions C), no acceptable decision rule can permit that agent to reject both bets.

A few notes are important. First: while some have questioned this constraint (see, for example, Moss (ms.) or Hedden (2013)), I will assume for the purposes of this paper that *Acceptance* is a requirement of rationality. Many of us have a strong pre-theoretic commitment that turning down a sure \$5 is irrational. Any decision rule that invalidates this intuition has failed to capture a basic normative commitment and ought to be rejected. Second: notice that our best theory for precise credences conforms to this requirement. For an agent that knows she is in conditions C, expected utility maximization never recommends that the agent reject both bets—regardless of the precise probability she assigns to q . Finally: note that the requirement is *not* that the agent accept both bets for the sure \$5. If Annie is especially confident in q or $\neg q$, it may be appropriate for her to reject one or the other bet, so long as she accepts at least one of the bets.

Permissive appears to violate *Acceptance*. According to some probability functions in Annie's representor, rejecting *Bet 1* maximizes expected utility. Thus, according to *Permissive*, it is permissible for Annie to reject *Bet 1*. According to other probability functions in her representor, rejecting *Bet 2* maximizes expected utility. Thus it is permissible for her to reject *Bet 2*. However, according to *Acceptance*, it is not permissible to reject both bets, so *Permissive* is apparently in violation of *Acceptance*.

2.3 Planning rules

The reader may complain about the above reasoning. On any plausible deontic logic (with ' P ' as the sentential operator 'Permissible'): $\{P(\phi), P(\psi)\} \not\vdash P(\phi \wedge \psi)$. While *Permissive* permits Annie to reject *Bet 1* and permits Annie to reject *Bet 2*, a permissive rule need not permit her to reject both bets. *Permissive* was designed to evaluate individual bets, but *Acceptance* is a constraint on evaluations of a diachronic series of bets. Perhaps a fan of imprecise credences can exploit this invalidity in order to adapt *Permissive* with rules for decision problems that include several choices across time. The following is a simple version of such a rule in which agents form plans for the sequences of decisions they will face and proceed to enact those plans:

Decision Rule: Planning A plan of actions is admissible just in case according to some credence function in the agent's representor that plan maximizes expected utility. An action α is permissible just in case the agent's past actions and α form part of an admissible plan.

According to *Planning*, Annie should evaluate plans of bets and then set off enacting one of these plans. For Annie, these admissible plans will never include the rejection of both bets. There are plans that allow her to reject or accept each bet; however the plan of rejecting both bets is inadmissible—for all credal functions in her representor, the expected utility of accepting both bets is greater than rejecting both bets. Supposing Annie rejects the first bet, rejecting the second bet would complete an inadmissible plan and is therefore impermissible according to *Planning*. Thus, *Acceptance* appears to be met.

Elga formulates a powerful argument against such proposals. Consider another imprecise agent, Sally, that is presented with *Bet 1* and *Bet 2* at times t_1 and t_2 respectively. (The bets are presented in conditions C, and Sally is certain of this set-up.) Suppose further that Sally has the same credences as Annie with respect to q . Compare the following two scenarios. Scenario 1: Sally has rejected *Bet 1*, and is faced with *Bet 2*. Scenario 2: Sally is only presented with *Bet 2* at time t_2 .

Elga argues that Sally's evaluation of *Bet 2* ought to be the same in the two scenarios. If Sally, after rejecting *Bet 1*, were to evaluate *Bet 2* differently than she would were she presented only with the second bet, then she would appear to be engaged in a type of 'sunk cost' reasoning, allowing irrelevant past decision to influence her present decision. Short of defending a bold new *sui generis* reason-giving mental state (like an 'intention'), Sally appears to have all the same reasons for action in the two scenarios. Thus, no rule can place different constraints on Sally in the two scenarios. Following McClennen (1990), let's call this final constraint *Separability*:

Requirement: Separability A decision rule must impose the same constraints on an agent who is presented with both bets and has rejected *Bet 1*, and an agent (with the same utility function and credence in q) that is presented with only *Bet 2*.

Several philosophers have questioned the requirement of *Separability*.⁴ Nevertheless, let's grant Elga this final constraint. Of course, *Planning* violates *Separability*. In Scenario 1, Sally has rejected the first bet and so must accept the second bet. In Scenario 2, the constraints are different; Sally is free to accept or reject the second bet.

Proposals linking Annie's decisions across time fail to meet *Separability*: *Planning* proposals bind together her evaluation of *Bet 2* with her evaluation of *Bet 1*, imposing different requirements for Sally in the two scenarios.⁵ *Separability* thus

⁴ See, most prominently, McClennen (1990).

⁵ Those familiar with the literature will note that the decision rule *Planning* above differs from two other diachronic decision rules discussed by Elga (2010): Narrowing and Sequence proposals. Narrowing proposals require agents to change their credal state after acting on the first bet. Such rules are antithetical to the evidentialist motivations of many imprecise Bayesians and are in violation of *Separability*. Sequence rules evaluate *sequences* of action for rationality directly; they include irreducible diachronic ought claims that take wide-scope over sequences of actions without issuing verdicts on the permissibility of individual actions. Sequence proposals raise a host of thorny questions. For my purposes here, I will

has the effect that Annie must evaluate her decision at t_2 independently of her decision at t_1 . This conclusion also provides Elga with the missing premise required to reject permissive rules: $(P(\varphi) \wedge P(\psi)) \implies P(\varphi \wedge \psi)$. Permissive rules that allow the agent to reject *Bet 1* or *Bet 2* must, by *Separability*, allow the agent to reject both *Bet 1* and *Bet 2* and are thus in violation of *Acceptance*.

At this point, Elga takes himself to have successfully partitioned the space of possible decision rules and systematically rejected all available options for the imprecise Bayesian. The complete set of norms—including both practical and epistemic norms—must be consistent. If no set of practical norms can consistently validate imprecise credences and Elga's three constraints, something has to go. Elga concludes that we should reject imprecise credences; others (Moss ms.; Hedden 2013; McClennen 1990) argue that we should reject one of Elga's constraints.

This is the stand-off that the literature has settled on.⁶ However, as we will make clear in the next section, an alternative strategy remains for supporters of imprecise credences.

3 A forward looking decision rule

In this section, I will argue that Elga is wrong in concluding that no decision rule for imprecise agents can simultaneously meet *Imprecision*, *Acceptance*, and *Separability*. In particular, I think that we can respect the intuition behind *Separability* without giving up on all diachronic rules. The decision rule that I will propose is inspired by two observations.

Observation 1. *Separability* prohibits Annie's decision on *Bet 2* from depending on her prior decision regarding *Bet 1*. However, there is nothing about *Separability* that prevents Annie, when considering *Bet 1*, to also consider her future decision regarding *Bet 2*. We need a decision rule that ties Annie's consideration of *Bet 1* with her future decision problem with respect to *Bet 2*, without tying Annie's consideration of *Bet 2* with her past decision problem with respect to *Bet 1*. In other words, we need Annie to act on a *forward looking* decision rule, in that the rule evaluates sequences of *present and future* actions, without concern for what has happened in the past.

As we'll see, this has implications for our ability to coordinate actions across time. *Separability* prevents agents from coordinating present action with past action—*Separability* prevents backwards-looking coordination. However, agents can coordinate present action with *future* action in a way that ensures diachronic consistency—*Separability* does not restrict forward-looking coordination.

Footnote 5 continued

take it for granted that such rules are off-the-table: a decision rule must issue in verdicts of the permissibility of an agent's decision at a particular time.

⁶ After drafting this paper, I've discovered other authors working on related projects. Chandler (forthcoming), Sahlin and Weirich (2013) and Bradley and Steele (ms.) each argue that, using backwards induction, the popular decision rule Γ -MaxiMin requires Elga's agent to accept both bets. (The rule, however, has well-known problems of its own and, most importantly, is not a 'permissive' rule which is a property favored by most imprecise Bayesians.) In their independent manuscript, Bradley and Steele also allude briefly to a rule similar to the one I explore below.

Keeping this first observation in mind, Annie can deploy the following intuitive reasoning when considering *Bet 1*:

Should I accept or reject *Bet 1*? Well that depends on what I do in the future with respect to *Bet 2*. I can either accept both *Bet 1* and *Bet 2*; I can accept *Bet 1* and reject *Bet 2*; or I can reject *Bet 1* and accept *Bet 2*. Between those sequences, I don't mind. Certainly, however, it would be bad if I reject both bets—that I'm sure about! But I also know that in the future, I won't care about my present decision: (by *Separability*) my future-self won't consider what I do now when confronted with *Bet 2*. Thus, she'll be undecided between accepting and rejecting *Bet 2*.

Notice the structure of Annie's reasoning above. When presented with *Bet 1*, Annie considers all plans or sequences of future actions and determines which of those plans are admissible. However, unlike *Planning*, she does not naively set off attempting to enact just any of these admissible plans. Instead, she pauses to consider the decision problem her future-self will face at t_2 , under the assumption that her future self, in accordance with *Separability*, won't consider past actions.

At this point, a second observation becomes apparent. Above, we suggested that Annie concern herself with the diachronic rationality of her present and future actions. (Picturesquely, Annie concerns herself with the time-worm that begins at the time of action.) Note further, however, that if Annie accepts the first bet, she will *guarantee* herself an admissible sequence of future actions, no matter what her future self chooses to do. In this way, she can coordinate her present action with her future actions to ensure the diachronic rationality of her future self. Annie can continue her reasoning from above:

Given that I can't control my future self's decision, the safest thing to do is to just accept *Bet 1*—that guarantees me an acceptable sequence of actions.

One way to spell out this reasoning is with the following decision rule, which I call *Forward Looking*. *Forward Looking* is a modification of *Planning*—it evaluates *sequences* of actions instead of treating each action independently. However, it incorporates our two observations from above: Annie considers the diachronic rationality of her future actions without concern for past actions and coordinates her present actions with her future actions by opting for an action that guarantees her a rational sequence of future actions. Following Weatherston (ms. a), I formulate the rule for agents faced with a choice between two bets.

3.1 Terminology

First, let's define some simple terminology. Suppose an agent at time i will be faced with a series B of choices between two bets. $B = (b_i, b'_i, b_{i+1}, b'_{i+1}, \dots, b_{i+n}, b'_{i+n})$ (where the subscript represents the time of each bet). At each time j , the agent can choose either bet b_j or b'_j . Call a possible sequence of actions s , and the set of all such sequences S . Thus, we have, $s \in S = \prod_{j=i}^{i+n} \{b_j, b'_j\}$.

A sequence, s , is *admissible* iff according to at least one probability function c in your representor, the expected utility of enacting s is greater than or equal to all other sequences $s' \in S$. Call the set of admissible sequences A .

Importantly, note that the decision problem that an agent faces changes after each move. Over time, the series of choices the agent faces (B) changes, the set of possible sequences (S) changes, and the set of admissible sequences (A) changes. So, we ought to index these to the time the agent is acting: B_i, S_i, A_i .

3.2 Forward looking

Above, we've described what constitutes an admissible sequence. A decision rule, however, must tell us which *actions* are permissible.

Forward Looking: For some agent at time i , that is certain she is in conditions C with respect to a finite series of choices between bets, B_i , when faced with a choice between b_i and b'_i , an action α_i is permissible just in case:

- i That action is part of some admissible sequence beginning at time i .⁷
- ii That action is not *sequence-dominated*. An action α_i is *sequence-dominated* iff
 - (a) for some sequence of future permissible acts, doing α_i in conjunction with that sequence is not an admissible sequence⁸; and
 - (b) there is some other action α'_i such that for all sequences of future permissible acts doing α'_i in conjunction with that sequence is an admissible sequence.⁹

3.3 Comments

Intuitively, part (i) of the rule states that for an action to be permissible, it must be validated by an admissible sequence of actions. The requirement, however, is importantly different from *Planning* in that only plans beginning at the time of action are considered; in this way, we encode our first observation from above in order to remain consistent with *Separability*. Part (ii) is meant to capture our second observation: if you can act in such a way that ensures your future-self will perform a sequence you now deem acceptable, you ought to choose that act. Whenever possible, this requirement prevents the agent from adopting an unenforceable plan—a plan that the agent knows may fail to ‘make sense’ to her future-self at the next round of the decision problem.

⁷ $\exists x, y, \dots$ such that $(\alpha_i, x, y, \dots) = s_i \in A_i$.

⁸ There is some sequence of future permissible acts, $(\alpha'_{i+1}, \alpha'_{i+2}, \dots)$ such that $(\alpha_i, \alpha'_{i+1}, \alpha'_{i+2}, \dots) \notin A_i$. Note the appeal to future permissible acts is not circular, as there will be some last action in the sequence of acts.

⁹ For all sequences s_{i+1} of future permissible acts, $(\alpha'_i) \triangleleft s_{i+1} = s_i \in A_i$.

Consider how the rule applies to Annie's decision problem. At time t_1 Annie is choosing between accepting Elga's first bet (b_1) or rejecting that bet (equivalently, accepting an alternative bet (b'_1) that pays \$0 in all outcomes). And she faces a similar choice at time t_2 (between b_2 and a bet b'_2 that pays \$0 in all outcomes). The set of admissible sequences, $A_1 = \{(b_1, b_2), (b_1, b'_2), (b'_1, b_2)\}$, includes both accepting the first bet (b_1) or rejecting that bet (b'_1). Both accepting or rejecting *Bet 1* are validated by some admissible sequence, and thus both actions satisfy requirement (i) of the rule. However, Annie does not naively set out enacting just any plan. Annie, as a sophisticated agent that reasons about her future decision problem, notices that rejecting *Bet 1* (b'_1) is *sequence-dominated* by accepting *Bet 1* (b_1).

To see this consider the decision problem that Annie will face at t_2 . At that stage in the decision problem, the set of admissible sequences include accepting *Bet 2* and rejecting *Bet 2*—there are credence functions that recommend either action. Thus we have $A_2 = \{(b_2), (b'_2)\}$. (The 'sequences' that begin at t_2 only include one bet, so neither action is sequence-dominated.) Thus, accepting *Bet 2* and rejecting *Bet 2* are both permissible actions for Annie at time t_2 . Remember, however, that Annie *knows this fact* at time t_1 . Given that both acts will be permissible for her future-self, if she rejects *Bet 1*, her future-self may complete an inadmissible sequence of actions by rejecting *Bet 2*. If, however, Annie (at t_1) accepts *Bet 1*, she guarantees herself a sequence of actions that she finds admissible at t_1 . By requirement (ii), Annie ought to accept the first bet.

Notice the following attractive features of *Forward Looking*. First, in the case of a single-bet sequence, *Forward Looking* reduces to *Permissive*. Because the 'sequence' consists of a single bet, no action is sequence-dominated. The only constraint that needs to be met for an act to be permissible is (i), the action needs to be part of some admissible sequence; when the sequence consists of a single bet, an action is permissible iff that action maximizes expected utility according to some probability function in the agent's representor. Imprecise Bayesians that find permissive choice rules attractive should be encouraged by this result. Second, if the agent's credence is precise, *Forward Looking* reduces to expected utility maximization. If there is only one credence function in the agent's representor (or all credence functions agree with respect to the relevant propositions), constraint (i) requires the agent to choose an action that maximizes expected utility.

Finally, notice that *Forward Looking* meets Elga's three constraints. *Separability* is met because the decision rule makes no reference to past decisions; the rule references *future* permissible actions to constrain present decisions, but does not reference past decision-making. *Imprecision* is met because in single-bet cases the decision rule reduces to *Permissive*. As we've explained above, *Acceptance* is met because Annie accepts the first bet.

4 Objections considered

The primary purpose of this paper was to demonstrate that there are decision rules that meet all three of Elga's requirements of rationality. If the previous section was successful, we have already demonstrated this result—even if the proposed decision

rule turns out to be objectionable for other reasons. Nevertheless, in this section I explore various potential objections to *Forward Looking* in order to suggest that *Forward Looking* is not obviously objectionable. In Sect. 5 I explore the limitations of the rule.

4.1 Ad hoc

One strategy for objecting to *Forward Looking* is to claim that the rule is not well-motivated. The requirement that an action not be sequence-dominated—formalized as requirement (ii)—is particularly susceptible to this challenge. Why, the objector may ask, should I be *forced* to accept *Bet 1*, given that another sequence of actions is also admissible? Avoiding sequence-dominated actions is unduly restrictive—and postulating such a requirement simply in order to solve Elga’s problem is *ad hoc*.

Response: avoidance of sequence-dominated actions is, in fact, a highly intuitive norm. Whenever possible, we should act now to ensure that we follow through on our plans. Note further that sequence-domination is only a weak, tie-breaking requirement for an imprecise agent that is undecided between various admissible sequences. If there is only one admissible sequence available to the agent, the requirement is not binding.

Moreover, the constraint has a well-motivated parallel in game theory. Several economists adopt the following normative constraint on rational strategies in game-theoretic contexts (paraphrased from Weatherson (ms. b, pp. 11–12)):

Elimination of Weakly Dominated Strategies: No permissible strategy is *weakly dominated*. A strategy s_1 weakly dominates strategy s_2 iff


- (a) for some combination of moves by other players (and states of the external world), playing s_1 provides a greater payoff than playing s_2 ; and
- (b) for all combination of moves by other players (and states of the external world), playing s_1 provides at least as high a payoff as playing s_2 .

Intuitively, *Elimination of Weakly Dominated Strategies*, states that it would be irrational to play a given strategy if some other strategy exists that guarantees you at least as good of an outcome, no matter what strategies the other players play, with the possibility of a better outcome, given some strategy that the other players may play. Take, for example, the game given by the decision tree in Fig. 1.

Player 1 has two available actions: A_1 or R_1 . Similarly, Player 2 has two available actions: A_2 or R_2 . Payouts for outcomes are given as ordered-pairs with Player 1’s payout listed first, and Player 2’s payout listed second. In this game, Player 1’s strategy of playing R_1 is weakly dominated by the strategy of playing A_1 : by playing A_1 , Player 1 will guarantee herself *at least* as good of an outcome as playing R_1 —and potentially a better outcome if Player 2 decides to play R_2 . While some have questioned *Elimination of Weakly Dominated Strategies*, many game-theorists have appealed to the norm. Those who accept this norm view playing R_1 as irrational.

Treating the temporal parts of Annie as players, is there a lesson to be drawn from the above game? Of course, the game does not *exactly* model the situation Annie is in. The game models two agents with different utility functions (whereas

Fig. 1 Example of weakly dominated strategy



Annie's utility function is stipulated not to change between bets) and no uncertainty is represented. However, if we allow agents to evaluate sequences of future actions as admissible or inadmissible, there is a parallel between the game above and Annie's situation.

Recall the rationale of *Forward Looking*. Agents form preferences over sequences of future actions, deeming some acceptable and some unacceptable. Note that Annie at t_1 prefers different sequences of actions to Annie at t_2 because they face different diachronic decision problems over which to plan. At t_2 , Annie doesn't concern herself with her past action. (Agent's don't plan for the past!) Sometimes, however, different acceptable sequences recommend performing different actions—the question of how to act remains underdetermined.

The modest suggestion is that, in such cases, the agent ought to avoid sequence-dominated actions. The motivation behind the elimination of sequence-dominated actions parallels that of *Elimination of Weakly Dominated Strategies*: rejecting the first bet is 'weakly dominated' by accepting the first bet, because accepting the first bet guarantees the agent a preferable sequence of actions no matter what her future-self chooses to do. This similarity is enough to rebut charges that elimination of sequence-dominated actions is *ad hoc*.

4.2 Requires fragile conditions

It may be pointed out that the conditions under which *Forward Looking* is applicable are quite fragile (recall the conditions C above). In cases that depart from these conditions, *Forward Looking* doesn't apply. For instance, suppose that Annie is presented with *Bet 1* and *Bet 2* separately. According to this modification, Annie does not believe at t_1 , that she will be presented with *Bet 2* at t_2 . In this scenario, *Forward Looking* evaluates each bet individually; Annie is allowed to reject both bets and give up a sure \$5.

In response to objections such as these, I need only point out that Elga cultivates the normative intuition behind *Acceptance* in comparatively fragile conditions. If the case is modified such that *Forward Looking* allows Annie to reject both bets, it becomes difficult to justify *Acceptance* as a requirement of rationality. Take for instance the modification discussed above: suppose that Annie is not told at t_1 that she will be presented with *Bet 2* at t_2 . In this case, I don't find it irrational for Annie to reject both bets. Of course, doing so would be sub-optimal, given that she could have accepted both bets for a sure \$5. Yet, this cost is the result of Annie not having complete information—of Annie not knowing the set-up of the scenario in advance. Here is another way to see the point: in the modified scenario above, Annie *objectively* ought not reject both bets. She could have done better by accepting both

bets—her ends would have better realized by another course of action.¹⁰ However, decision rules are norms of rationality. They encode the requirements of the *subjective* ought—a sense of ought that respects our limited evidential state and is tied to our practices of praise and blame. If Annie does not know in advance that she will be presented with both bets, it is not clear that she subjectively ought not reject both bets. She has done nothing *irrational*: it is not clear that she would be *blameworthy* for her actions in the way that she would be had she known about the set-up in advance. At the very least, Elga would require much more intuition-pumping in order to argue that *Acceptance* is a binding constraint in modifications of conditions C.

4.3 Consequentialism

A final strategy for objecting to *Forward Looking* generalizes some comments Elga makes in the context of explaining our commitment to *Separability* and wields those comments against *Forward Looking*.¹¹ Recall that our intuitions in Elga's Sally case militated in favor of *Separability*. Elga's attempts to explain this result with the following appeal: "...in each situation, the consequences of accepting [Bet 2] and rejecting [Bet 2] are exactly the same in every respect that Sally cares about...So it can't be that rationality imposes different requirements on her in the two situations." (p. 8). Elga seems to be subscribing to the following strict consequentialist principle ('SCP'): *if the consequences of an action, in terms of utility, are the same in two scenarios, then a decision rule must not distinguish between the two scenarios.*

Forward Looking, while it conforms to *Separability*, is in violation of SCP. Consider the following variant of Annie's decision problem.¹² Annie has a twin, Anjali, who is in the same credal state as Annie. However, Anjali is not a miser. Instead, her utility function increases linearly with the dollars that Annie receives. Annie will be presented with *Bet 1* and, afterwards, Anjali will be presented with *Bet 2*. Although Anjali will decide with respect to *Bet 2*, the financial outcome of the bet will be borne by Annie.

What should Annie do when presented with *Bet 1*? Because she will not act with respect to *Bet 2*, *Forward Looking* allows her to accept or reject the bet. However, this is in violation of consequentialism—the financial consequences of Annie's decision are the same in the original scenario and in the twin-variant. Why should the fact that it is Annie—as opposed to Anjali—acting at t_2 make any difference to Annie's action at t_1 when the consequences of her action are the same? And, by stipulation, Annie only cares about money; she doesn't care about, for instance, the identity of the person presented with *Bet 2*.

¹⁰ Indeed, even accepting both bets is not what she objectively ought to have done—depending on whether p is true or false, she would have maximized her ends by accepting either *Bet 1* or *Bet 2* (but not both).

¹¹ Thanks to Sarah Moss and Brett Topey for raising this objection.

¹² Thanks to Tom Dougherty for raising this example and Miriam Schoenfield for pushing me on it.

I will not attempt to adjudicate what it takes for a decision rule to count as ‘consequentialist’. My purpose, again, is not to defend *Forward Looking* as the correct decision rule for imprecise agents. Instead, I am attempting a more modest goal; I merely want to convince the reader that conceptual space exists for a decision rule that abides by *Separability* and meets Elga’s other two constraints. It suffices, then, to cast some doubt on Elga’s appeal to SCP as the principle that explains our commitment to *Separability*.

First note that, as stated above, SCP is simply false. Classical, precise, expected-utility maximization *already* violates SCP. According to the expected utility maximizer, it is not the case that an action is permissible just in case it maximizes utility—the action must maximize *expected* utility. The agent’s beliefs matter just as much as the consequences of the action. If the consequentialist principle is to remain plausible, it must be further tailored (‘SCP*’): *if the expectation of utility of an action is the same in two scenarios, then a decision rule must not distinguish between the two scenarios*. *Forward Looking* violates SCP*.

Once we state the principle in this way, it becomes clear that appealing to the sort of consequentialism that would serve as a basis of criticism for *Forward Looking* simply begs the question against theorists that are investigating decision rules which aren’t merely forms of expected utility maximization. The complaint amounts to an accusation that *Forward Looking* considers factors extraneous to the expected-utility of an action, such as the structure of the agent’s diachronic decision problem. And that’s an accusation that shouldn’t trouble a defender of *Forward Looking*. (Compare an accusation that the expected-utility maximizer violates SCP by considering extraneous factors like the agent’s *beliefs*.)

My objector may retort: ‘I am not begging the question against you, for I am not simply asserting SCP* without support. My consequentialist principle is the only way to explain our intuitions behind *Separability*. Its ability to explain those intuitions is evidence for the principle. And if your rule is blatantly inconsistent with the principle, then so much the worse for your decision rule.’

Reply: our judgments of rationality do seem attuned to factors that go beyond the consequences of our actions, including facts about personal identity. Note that, unlike in the original Annie case, there is no intuitive support for *Acceptance* in the twin-variant case. Our shift in intuitions between the case is *prima facie* evidence that facts of personal identity are normatively relevant. Similar remarks apply in precise contexts. Diachronic epistemic rules like conditionalization and reflection make reference to facts of personal identity. And the intuitions of irrationality that undergird the diachronic Dutch Book arguments that support these rules rely on the fact that the same agent is acting across time.¹³

Moreover, the gap between *Separability* and SCP* is larger than my opponent is letting on. There are other ways to explain our commitment to *Separability* that don’t entail SCP*. The issue may turn on more subtle issues of agency. For instance, rules that violate *Separability* may be seen as psychologically or motivationally

¹³ See Hedden (ms., Sect. 4) for a discussion of this feature of diachronic Dutch Books. (As a ‘time-slice rationalist’, however, Hedden takes this observation to undercut support for diachronic Dutch Book arguments.)

unfeasible.¹⁴ Such rules are naturally packaged with a commitment to a psychological attitude—an intention—that motivates an agent to follow through on the plans that they prefer. One possible reason for endorsing *Separability*, then, would be the denial that such attitudes can motivate an agent in this particular way.

Other potential motivations for *Separability* appeal directly to distinctions between the normative relationships that our present-selves bear to our past- and future-selves. Such distinctions apparently exist; my present-self bears a number of normatively significant relationships with my future-selves that are not borne by my past-selves. For instance, I feel fear for my future-selves in a way that I don't feel for my past-selves. And I exert causal influence on the future, but not on the past. As Ismael (2012) discusses, from our perspective as decision-makers, the future seems in our control in a way that the past does not. This seems evident in the way in which we plan; we consider plans for our future-selves without planning for our past-selves. Given the asymmetries in our normative relationships to the past and future, it wouldn't be surprising were our present-selves to have a responsibility for ensuring the diachronic consistency of their present and future actions, without being similarly constrained by their past actions.

Admittedly, before *Forward Looking* can be adopted in earnest, the remarks above would need to be fleshed out in much more detail. At this point, I've merely sketched some possible avenues for resisting the particular brand of consequentialism that would conflict with *Forward Looking*.¹⁵ Indeed, as I argue in the next section, there are simpler ways to strengthen Elga's original argument in response to this strategy of forward-coordination.

5 A supercharged good book

Forward Looking gives the imprecise Bayesian a well-motivated and highly natural rationale for avoiding Elga's diachronic book of bets. By accepting the first bet, an agent can coordinate her present actions with her future actions in such a way that ensures that her future-self will complete an admissible sequence. A natural question arises: will this strategy work in all diachronic-coordination puzzles? Is *Forward Looking* the answer to the imprecise Bayesians' prayer? Unfortunately, we will see that it is difficult to generalize our strategy of forward-looking coordination in scenarios involving more complicated books of bets.

Before exploring these more-involved diachronic puzzles, it is interesting to note that *Forward Looking* is a impressively robust decision rule; it appears to be

¹⁴ This is a complaint levied against Resolute Choice theorists. See Kavka (1983) for an initial discussion.

¹⁵ A route may exist between *Forward Looking* and a *sort of* consequentialist justification. Although unacceptable sequences cannot always be avoided (see Sect. 5), the consequences for an agent that avoids sequence-dominated actions are typically better than an agent that does not, in that the former often avoids completing an unacceptable sequence. So, there may be a rule-consequentialist justification in the offing for a defender of *Forward Looking*. (Thanks to Alex Guerrero for discussion here. See also Buchak (ms., Chapter 6) for more on consequentialist justifications and dynamic choice theories.)

immune to diachronic puzzles in a large class of decision problems. Trouble arose for decision rules that allowed agents to complete some sequence of actions while another sequence of action accorded the agent more expected-utility according to all members of the agent’s credal representor. When presented with a ‘good book’ of bets, the agent opted for a strictly worse book of bets.

Faced with simple ‘good books’ like Elga’s, in which a miser is presented with a finite series of bets on some proposition, each of which the agent can accept or reject, *Forward Looking* will *never* allow the agent to pass up a good book of bets for a strictly worse book of bets.¹⁶

Trouble re-arises for the imprecise Bayesian when we depart from these simple decision scenarios, and consider more complicated diachronic puzzles. One way in which to complicate matters is to present the agent with choices between three or more bets at a time. When we try to generalize the strategy of forward-looking coordination in the natural way, we will find that agents are unable to coordinate their present action with their future-selves in a way that guarantees an admissible sequence of actions is enacted. Consider the following game show in which an agent is presented with a good-book that is ‘supercharged’ to include three bets at each choice point.

Michael is in a boring game show that works in the following way:

- At $t = 1$, he will be given a choice of selecting Door 1, Door 2, or Door 3.
- At $t = 2$ he will be given a choice of selecting Door 4, Door 5, or Door 6.
- At $t = 3$ the selected doors will be opened.

Michael (who is certain that he will be in conditions C) is told that the payoffs behind the doors correspond to some proposition q and are as follows:

- Door 1: if q , \$9 will be behind the door; if $\neg q$, \$5 will be behind the door.
- Door 2: if q , \$5 will be behind the door; if $\neg q$, \$10 will be behind the door.
- Door 3: if q , \$8 will be behind the door; if $\neg q$, \$7 will be behind the door.
- Door 4: if q , \$4 will be behind the door; if $\neg q$, a penalty of \$4 will be behind the door.
- Door 5: if q , a penalty of \$4 will be behind the door; if $\neg q$, \$4 will be behind the door.
- Door 6: if q , \$2 will be behind the door; if $\neg q$, \$2 will be behind the door.

At time t_1 , Michael evaluates 9 possible sequences of actions given in the following table:

(if q , if $\neg q$)	Door 1	Door 2	Door 3
Door 4	(13, 1)	(9, 6)	(12, 3)
Door 5	(5, 9)	(1, 14)	(4, 11)
Door 6	(11, 7)	(7, 12)	(10, 9)

¹⁶ More carefully, I prove the following claim in the Appendix: Presented with a series $B = (b_0, b'_0, b_1, b'_1, \dots, b_n, b'_n)$ of choices between bets on the truth of some proposition q over time, a miser (who is certain she is in conditions C) that follows *Forward Looking* will always enact a sequence of actions that is admissible at $t = 0$.

The cells of the table indicate the payoffs, in dollars, for each sequence of actions. At t_2 Michael will evaluate three possible ‘sequences’ of single actions, given in the following table.

(if q , if $\neg q$)	
Door 4	(4, -4)
Door 5	(-4, 4)
Door 6	(2, 2)

Suppose that Michael is fully imprecise with respect to q ($c'(q) = [0, 1]$). Note that some sequences of actions at t_1 are strictly worse than others. For instance, according to all members of Michael’s representor, choosing Door 2 and Door 6 will provide Michael greater expected utility than choosing Door 1 and Door 5. Other similarly universally dispreferred sequences of actions are shaded in the above table. Note further that there is no universally dispreferred sequence of actions at t_2 : each action is recommended by some member of Michael’s credal representor.¹⁷

What doors should Michael choose in this decision-scenario? Strictly speaking, of course, *Forward Looking* is silent about such cases: the rule is only written for cases in which an agent is choosing between two bets at a time. However, it is not clear how Michael, at t_1 , could coordinate his present actions with his future actions in order to guarantee that his future-self will avoid an inadmissible sequence of actions. Given that it is admissible for Michael to choose any door at t_2 , no matter what door he chooses at t_1 , Michael risks his future-self completing an inadmissible sequence of actions.¹⁸

6 Conclusion

Suppose the reader by-and-large agrees with the discussion up to this point. She acknowledges the existence of a decision rule that avoids Elga’s diachronic puzzle, and delivers intuitive results when faced with simple good-book arguments. She also acknowledges the limitations of such a rule in more complicated diachronic decision problems. What is the ultimate upshot of these observations?

¹⁷ $c(q) = 1$ recommends opening Door 4; $c(q) = 0$ recommends opening Door 5; $c(q) = .5$ recommends opening Door 6

¹⁸ Although I don’t explore them here, there are plenty of other ways in which to complicate Annie’s decision problem so that the limited result of the Appendix does not apply. For instance: the bets could be on different propositions, Annie’s action with respect to past choice point could affect the utility associated with future bets, or the choices Annie makes in the past could affect which choices she faces in the future. As an instance of the last, consider an agent with two bets b_1 and b_2 that are incomparable in the sense that each maximizes expected utility according to some member of the agent’s representor. At t_1 , the agent must choose to sweeten exactly one of b_1 or b_2 by adding a small bonus prize to the bet – although the bets remain incomparable after sweetening. At t_2 , the agent chooses between the newly sweetened bet and the alternative unsweetened bet. Here it is difficult to see how forward coordination could prevent the agent from choosing an unsweetened bet. (Many thanks to Seamus Bradley for pointing out this example to me. A similar puzzle for imprecise preferences appears in Hedden (2013, Sect. 2.4)).

The lesson the reader ought to draw from the preceding sections depends, I think, on her antecedent beliefs about Elga's original argument. If, prior to reading this paper, the reader thought Elga's argument was effective in showing that imprecise credences are never ideally rational, then the upshot of the paper is as follows. We began by exposing a weakness in Elga's original argument, explored the extent to which that weakness can be exploited, and ultimately strengthened the argument by repairing his diachronic puzzle in the face of this exposed weakness.

Some readers, however, may have found Elga's argument unconvincing even *before* learning of the weakness discussed in this paper. In particular, some may have agreed with Elga's three constraints on decision making with imprecise credences, yet thought that the *pragmatic* difficulties of decision-making that Elga points to fail to warrant the *epistemic* conclusion that he seeks.¹⁹

From the perspective of such a reader, who is already pessimistic about the ability to 'depragmatize' Elga's good-book argument, the lesson of this paper may be broader. If we accept that ideally rational agents can hold imprecise credences, these agents need a decision theory! And *Forward Looking* may get us *part* of the way towards finding such a decision theory. The search for decision theories in unfamiliar contexts often begin with debates over the plausibility of necessary constraints for permissibility. For example, the literature on rational decision-making in game-theoretic scenarios is largely focused on constraints like *Elimination of Weakly Dominated Strategies*. While the results of Sect. 5 may lead us to conclude that *Forward Looking* is not an adequate decision rule on its own, the avoidance of sequence-dominated actions may at least constitute a necessary constraint on decision making with imprecise credences—one that delivers the intuitive results in a range of cases.

So: is the avoidance of sequence-dominated acts a necessary condition for the permissibility of acting with imprecise credences? On this question, I want to stay non-committal. Much more would have to be said to warrant a conclusive judgment. How defensible is this picture in which agents re-evaluate sequences of future actions without regard for the plans of their past selves? Should the rule be adapted for cases in which the agent has foreknowledge of her future actions?²⁰ Does the constraint generate unintuitive results and how should these results be weighed

¹⁹ Weisberg (forthcoming) presents an argument to this end. In the case of precise agents, susceptibility to Dutch Books is thought to demonstrate that the agent is committed to regard as fair a series of bets that are clearly not fair. Because our credences are thought to fix our commitments of fairness of bets, the inconsistency in our commitments—brought to the fore by Dutch Book susceptibility—is emblematic of an inconsistency of our credences. Weisberg argues that, in the case of imprecise agents, this last step of this argument fails. Roughly, he argues that imprecise credences often underdetermine commitments to the fairness of bets—imprecise agents may be undecided as to whether or not a particular bet is fair. The choices of an imprecise agent, then, do not correspond with commitments to the fairness of bets.

²⁰ Consider some contrived cases in which Annie is absolutely certain about her action with respect to Bet 2. If she is certain she will reject, it may seem strange to treat the case differently than a case in which she is presented with only the first bet. If she is certain that she will accept, it may seem unduly harsh to require Annie to accept the first bet. (Thanks to Carrie Ichikawa Jenkins and Jeff Russell for discussion here.)

against the intuitive results it successfully delivers?²¹ Challenges plainly remain for a would-be defender of the proposed constraint.

Nevertheless, I hope to have said enough to reinvigorate the dialectic by offering fans of imprecise credences a novel approach in their search for an adequate decision rule. Forward looking coordination, implemented through norms like sequence-domination, gives the imprecise Bayesian a powerful tool for balancing permissive decision-making with diachronic rationality.

Acknowledgments This paper has benefited greatly from conversations and comments from many friends and teachers. Thanks to Harjit Bhogal, Rachel Briggs, Lara Buchak, Tom Dougherty, James Joyce, Jason Konek, Leon Leontyev, Miquel Miralbes Del Pino, Miriam Schoenfield, Daniel Singer, Eric Swanson, and Brett Toppey. Special thanks to Seamus Bradley, Dmitri Gallow, Carrie Ichikawa Jenkins, Joshua Schechter, Sarah Moss, and Brian Weatherson. I am also grateful to audiences at the 2012 Australasian Association of Philosophy Conference, the 2012 Brown University Shapiro Graduate Philosophy Conference, the University of Pennsylvania and, especially, to participants of the Fifth Formal Epistemology Festival and the 2013 Bellingham Summer Philosophy Conference.

Appendix: proof

In this Appendix, I show that *Forward Looking* cannot be subjected to simple diachronic puzzles. More carefully, we show:

Main Claim: Presented with a series $B = (b_0, b'_0, b_1, b'_1, \dots, b_n, b'_n)$ of choices between bets on the truth of some proposition q over time, a miser (who is certain she is in conditions C) that follows *Forward Looking* will always enact a sequence of actions that is admissible at $t = 0$.

Proof Suppose the agent's representor is $c^r(q) = [\lambda, \eta]$. Call the probability functions that are members of the agent's representor her 'credal members'. Note that each choice at i between b_i or b'_i is associated with a 'pivot point' that bifurcates the unit interval such that one bet b_i maximizes expected utility according to all the credal members that assign probability to q greater than that pivot point. The alternative bet maximizes expected utility according to all members that assign probability less than that pivot point. Either action maximizes expected utility according to the member that assigns probability equal to that pivot point (Fig. 2). Let θ_i denote this member of the agent's representor that assigns probability equal to the pivot point at i .²²

Either (a) the pivot point for no choice points are in the agent's representor or (b) there is some *last* choice point, l such that the pivot point for that bet is in the agent's representor $\theta_l \in [\lambda, \eta]$. If (a) the result is trivial: all members of the agent's representor recommend the same action for each choice point, so that sequence of

²¹ For instance: one result of the rule is that the order in which the bets are presented to Annie affects which bets Annie must accept. Some have found this oddity objectionable (although I don't). (Thanks to Joshua Schechter for discussion here.)

²² When I talk of θ_i as a real number instead of as a probability function, I should be understood as referring to the pivot point at i (i.e. θ_i evaluated for the relevant proposition q).




Fig. 2 Pivot point of Bet i




Fig. 3 Case (iii)

actions will be enacted and is admissible. If (b), we show the following by induction:

Lemma: Any action that is permissible maximizes expected utility according to θ_l .

Base Case For any choice point $i \geq l$, any action that is permissible maximizes expected utility according to θ_i .

With respect to choice l , both actions maximize expected utility according to θ_l . With respect to choices i after l , by assumption $\theta_i \notin [\lambda, \eta]$. Thus, for each of these choices, there is a unique permissible action and the sequence of those actions maximizes expected utility according to all members of the agent’s representor. Thus, that sequence maximizes expected utility according to θ_l .

Induction Step Suppose for choices $i > m$, any action that is permissible maximizes expected utility according to θ_i . Then, for choice m any action that is permissible maximizes expected utility according to θ_l .

There are three cases: either (i) $\theta_m \notin [\lambda, \eta]$ (ii) $\theta_m = \theta_l$ or (iii) $\theta_m \in [\lambda, \eta]$ and $\theta_m \neq \theta_l$. Cases (i) and (ii) are trivial. If $\theta_m \notin [\lambda, \eta]$ then all members of the agent’s representor recommend the same action with respect to choice m , so that unique permissible action is recommended according to θ_l . If $\theta_m = \theta_l$ then both actions with respect to choice m maximize expected utility according to θ_m and thus both actions maximize expected utility according to θ_l .

In case (iii), some members of the agent's representor recommend the same action that maximizes expected utility according to θ_l ; some recommend the action that does not maximize expected utility according to θ_l (Fig. 3). However, by the induction hypothesis, the action that maximizes expected utility according to θ_l guarantees the agent an admissible sequence of actions (because all of the future actions will maximize expected utility according to $\theta_l \in [\lambda, \eta]$). And the alternative action can be combined with one of the permissible actions at choice l to form an inadmissible sequence. Thus the action that maximizes expected utility according to θ_l sequence-dominates the alternative action.

This proves the **Lemma**. The **Main Claim** follows immediately: all permissible actions maximize expected utility according to θ_l , thus the sequence of actions that is enacted is admissible.

References

- Bradley, S., & Steele, K. (ms.). Subjective probabilities need not be sharp.
- Buchak, L. (forthcoming). *Risk and rationality*. Oxford University Press.
- Chandler, J. (forthcoming). Subjective probabilities need not be sharp. *Erkenntnis*.
- Dorr, C. (2010). The eternal coin: A puzzle about self-locating conditional credence. *Philosophical Perspectives*, 24(1), 189–205.
- Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers' Imprint*, 10(5), 1–11.
- Hedden, B. (ms.). Time-slice rationality.
- Hedden, B. (2013). Options and diachronic tragedy. *Philosophy and Phenomenological Research*. doi:10.1111/phpr.12048.
- Ismael, J. (2012). Decision and the open future. In Adrian B. (Ed.), *The future of the philosophy of time* (pp. 149–168). Oxford: Oxford University Press.
- Joyce, J. (2011). A defense of imprecise credences in inference and decision making. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 4). Oxford: Oxford University Press.
- Kavka, G. S. (1983). The toxin puzzle. *Analysis*, 43(1), 33–36.
- McClennen, E. F. (1990). *Rationality and dynamic choice*. Cambridge: Cambridge University Press
- Moss, S. (ms.). Credal dilemmas.
- Sahlin, N.-E., & Weirich, P. (2013). Unsharp sharpness. *Theoria*. doi:10.1111/theo.12025.
- Weatherston, B. (ms. a). Decision making with imprecise probabilities.
- Weatherston, B. (ms. b). Lecture notes on game theory.
- Weisberg, J. (forthcoming). You've come a long way, Bayesians. *Journal of Philosophical Logic*.