Episteme, 11, 3 (2014) 277–289 © Cambridge University Press
doi:10.1017/epi.2014.8

SHOULD SUBJECTIVE PROBABILITIES BE SHARP?

SEAMUS BRADLEY AND KATIE STEELE
seamus.bradley@lmu.de and k.steele@lse.ac.uk

ABSTRACT

There has been much recent interest in *imprecise probabilities*, models of belief that allow unsharp or fuzzy credence. There have also been some influential criticisms of this position. Here we argue, chiefly against Elga (2010), that subjective probabilities *need not* be sharp. The key question is whether the imprecise probabilist can make reasonable sequences of decisions. We argue that she can. We outline Elga's argument and clarify the assumptions he makes and the principles of rationality he is implicitly committed to. We argue that these assumptions are too strong and that rational imprecise choice is possible in the absence of these overly strong conditions.

1. INTRODUCTION

That epistemic agents have partial beliefs is not very controversial. What structural rules these partial beliefs obey is a matter of quite some discussion. A standard way to argue for a particular structural requirement on partial belief is to demonstrate that to violate the requirement leads to some kind of irrationality. One important form of such an argument is the 'Dutch book argument' where the form of irrationality demonstrated is to accept an obviously bad set of bets, and the epistemic conclusion is *probabilism*. Thus there is a strong link between epistemology and decision theory.

Recently there has been a resurgence of interest in the question as to whether partial beliefs ought to be sharp: whether or not imprecise credence is permissible. The standard precise probabilist or *Bayesian* approach is to understand partial beliefs as conforming to the axioms of probability theory. This requires that every proposition be associated with a particular real number that represents the degree to which that proposition is believed. Some have wondered whether such sharpness is rationally required. Those who find that sharpness is unnecessary are said to subscribe to *imprecise probabilism*; otherwise known as indeterminate probabilism, unsharp credence or mushy credence. Standardly, these credences are represented by interval-valued functions, set-valued functions or sets of real-valued functions. We prefer to think of imprecise probabilities as sets of probability functions, and we call such sets *representors* (van Fraassen 1990).

Adam Elga's recent and influential paper *Subjective Probabilities should be Sharp* (Elga 2010) argues that the imprecise probabilist makes unavoidably bad sequences of decisions, and thus, that subjective probabilities should not be imprecise. That is, Elga argues that under certain conditions, imprecise credences lead to bad decisions and so are irrational. He therefore subscribes to the aforementioned tradition that practical consequences of your beliefs are relevant to their rationality.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

Given the tight connection between epistemology and the practical consequences of belief, we call the package of your belief state and choice rule your *decision theory*. We broadly agree with Elga's method: decision theories must answer to criteria of rationality that have to do with decisions, including sequences of decisions. Bets are paradigm cases of decisions, so decision theories must answer to criteria of rationality that have to do with (sequences of) bets. We disagree, however, with important details of Elga's analysis. His constraints on rational sequences of decisions are not plausible and thus his conclusions are wrong. The structure of the paper is as follows: we outline Elga's series of bets. Then we point out the strong assumptions he relies on to draw his conclusions. We then discuss the right way to approach sequential decisions, and indicate how our criteria of reasonableness differ from Elga's. We further consider the performance of two specific imprecise choice rules with respect to Elga's sequence of bets.

Before proceeding, a couple of basic points about imprecise decision theories: in this context, your degree of belief in a particular proposition is represented by the set of values assigned to it by the set of functions that represent your belief state.[1] Such a representation of belief also then gives rise to an *imprecise expectation*. That is, you assess the value of a possible act by looking at the probability-weighted values or utilities of its outcomes. If the probability is imprecise, then so too will be the utility. Every probability function in your representor gives rise to an expected utility. The imprecise expectation can be represented by the set of these utilities. The best act, the act you should choose, is the one with the highest expected utility. But if the expectations are imprecise, it can happen that the intervals of expected utility 'overlap'. It then becomes much harder to say which act is the best. The acts may be regarded as *incommensurable*. We will see an example of this in the next section. Incommensurability is vital for Elga's argument.

2. A GREAT SERIES OF BETS

Let us set the scene. Elga outlines a decision problem that looks like this:

> I'm going to offer you a great series of bets on H (where H is some particular proposition):
> **Bet *A*** If H is true, you lose \$10. Otherwise you win \$15.
> **Bet B** If H is true, you win \$15. Otherwise you lose \$10.
> First I'm going to offer you Bet A. Immediately after you decide whether to accept Bet A, I'm going to offer you Bet B. (Elga 2010: 4).

Judged as a four-way choice between 'take both bets', 'take bet A', 'take bet B' and 'take neither bet', it should be obvious that taking neither bet is unwise. If you have high credence in H, then bet B looks good; low credence in H makes A look better. Middling credence in H suggests taking both bets: that guarantees you a payout of \$5. Put another way, whatever your belief about H, you shouldn't refuse both bets: refusing both bets is *dominated* by accepting both. An act or option X is *dominated* by another option Y when the expected utility of Y is at least as great as the expected utility of X for all probability functions in your representor, and for at least one probability function in your representor,

1 There are reasons to be unhappy with this identification, but they need not concern us here.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

the expected utility of *Y* is strictly greater than the expected utility of *X*.[2] Taking both bets guarantees you \$5 whatever you believe about *H*, so refusing both (net gain \$0) is dominated. The assumption here – and throughout the paper – is that utility increases linearly with money. Everyone agrees that choosing dominated options is irrational.

In what follows we will use the following shorthand when discussing the various options: *AB* means you take both bets; *AN* means you take bet A but refuse bet B; *NB* means you refuse bet A but take bet B; and *NN* means you refuse both bets. If all these combinations are *available* to the agent, the choice of *NN* is irrational, since *AB* is *universally* better than it; better whatever you happen to believe about *H*. Now, one of *NB* or *AN* might be better still, depending on your beliefs, but the main point is that *NN* is dominated by *AB*. Every possible probabilistic belief function ranks *AB* above *NN*: every precisification of your imprecise credal state ranks *AB* above *NN*.

Imagine that you have no evidence whatsoever about what kind of proposition *H* might be. So, following the imprecise credence model, it is plausible that your representor contains distributions that assign *all possible probability values* to the event *H*. That is, the set of values the distributions in your representor assign to *H* covers all of the unit interval [0, 1].[3] In this case the set of expectations assigned to Bet *A* can be represented as [−10, 15] and likewise for Bet B. Refusing either bet has an expectation of 0. The expectations for the compound bets *AB*, *AN*, *NB* and *NN* are represented as follows:

- *AB* {5}
- *AN* [−10, 15]
- *NB* [−10, 15]
- *NN* {0}

These expectations underscore the point above that *NN* is dominated by *AB*. For this particular agent with beliefs as specified, *AB*, *AN* and *NB* are all incommensurable (roughly, they have 'overlapping' sets of expectations). In fact, there are some probabilities in your representor that rank *AN* as best (highest expected utility), some that rank *NB* as best, and some that rank *AB* as best. But no probability in your representor ranks *NN* as best. In a four-way choice, any sensible imprecise decision rule should satisfy this dominance reasoning, and therefore rule out *NN* as an admissible choice.

Elga is not interested in the four-way choice as in Figure 1, but in *sequential* choices: where you choose whether to take bet *A* first, and then decide whether to take bet *B*. So what goes wrong in the imprecise context? We now go through Elga's analysis of how imprecise choice rules fail in this kind of example. Take a very simple imprecise choice rule which says that any option is an admissible choice,[4] unless some other option

2 This definition pertains to *expected utility dominance*. There is also a logically stronger definition of dominance: *X* is state-wise dominated by *Y* when *Y* leads to at least as good an outcome (i.e. at least as high utility) as *X* in every possible state of the world. For Elga's decision problem, the package of both bets state-wise dominates taking neither bet, as well as expected utility dominating it. This is not always the case, and the definition of dominance we use in the text will always be expected utility dominance.

3 This is not exactly the belief representor that Elga assumes for his example, but the difference does not matter for the analysis.

4 Imprecise choice rules are standardly defined in terms of what options they permit as *admissible* choices from amongst a set of available options. We use this terminology in the remainder of the paper. Note

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

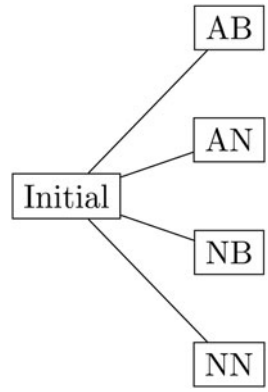


Fig. 1. A four-way choice

dominates it. In the first choice, between *A* and *N*, neither option dominates. So in particular, *N* is an admissible option. Now in the second choice, again, neither of *B* or *N* dominates. So again, *N* is an admissible option. An agent who reasons like this may end up having chosen *NN*, or rather, it is permissible for this agent to end up choosing *NN*.

Elga's strategy is to show how each of a number of proposals for imprecise choice fail: either by permitting the choice of *NN* in his sequential betting example, or by having some other defect. Thus no imprecise choice rule is acceptable. From this Elga argues that no imprecise credence can be rational and thus that subjective probabilities should be sharp. We will focus on the former of Elga's criteria for acceptable imprecise choice rules – that they do not permit the sequence *NN* in his betting example – because it is this criterion which rules out the choice rules we regard as most plausible. We take issue with Elga's analysis of which imprecise choice rules *do* permit the sequence *NN*, and we furthermore take issue with even using 'does not permit the sequence *NN*' as a criterion of rationality.

In arguing for his conclusions, Elga makes use of some assumptions that we find objectionable. These assumptions are clearest in his discussion of a proposed approach to sequential rationality that he calls 'SEQUENCE'. The SEQUENCE proposal is as follows:

> Just as individual actions can be assessed for rationality, so too can sequences of actions. And it can happen that a sequence of actions is irrational even if each of its elements is rational. In particular, suppose that an agent has rejected both bets in the Bet A/Bet B situation. Then her first action – rejecting Bet A – was rationally permissible. And her second action – rejecting Bet B – was rationally permissible. But her performing the *sequence* of actions 'reject-Bet-A-then-reject-Bet-B' was rationally impermissible. (Elga 2010: 9)

We find the SEQUENCE proposal misdirected, as may be inferred from our discussion in Section 4, but that is not Elga's criticism. Here is Elga's argument against SEQUENCE:

> Consider … two situations in which Sally is considering Bet B. In the first situation, she has previously rejected Bet *A*. In the second, she was never offered Bet *A* at all. In the two situations, Sally

also that criteria of rationality for imprecise choice rules are often stated in terms of what options a choice rule *should (not) make admissible in a given context*.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

> faces choices that are exactly the same in every respect she cares about ...
>
> So it must be that rationality imposes the same constraints on her in the two situations. (Elga 2010: 9)

Elga thinks SEQUENCE conflicts with the above in that it allows a choice rule to make different demands of an imprecise agent in situations where, according to Elga, the demands of a choice rule should be the same. He thus concludes that SEQUENCE is inconsistent. We are using what Elga says about SEQUENCE to highlight some assumptions he makes that we consider misguided.

3. ELGA'S OBJECTIONABLE ASSUMPTIONS

To set the stage for our positive proposal let us specify the set of assumptions that Elga seems implicitly committed to, given his discussion of the SEQUENCE proposal:

1. *Retrospective Rationality*: An important assessment of rationality (of an agent, or rather their decision theory) concerns the sequence of decisions the agent ends up making. If the sequence of decisions that an agent ends up making is dominated by another sequence of decisions the agent had the opportunity to make, then the agent's sequence of decisions is irrational, from this retrospective perspective.
2. *At-a-time Rationality*: Individual decisions should be made in accordance with what is rational for the agent *at that time*.
3. *Criticism of* SEQUENCE: It should never be the case that an agent abides by *At-a-time Rationality* when making decisions but ends up being *Retrospectively Irrational*.

These three assumptions, as interpreted by Elga, entail a conclusion that we refer to as the *Strong Package Principle*:

> *Strong Package Principle*: Individual decisions should be rational in isolation (respecting *At-a-time Rationality*) and should also straightforwardly agglomerate, such that the resulting sequence of decisions is *Retrospectively Rational*.

We now discuss our disagreements with Elga regarding the above assumptions. We effectively deny assumption 1 (and thus 3) and we have a different interpretation of assumption 2. We will then discuss the conclusion that follows from these assumptions: the Strong Package Principle.

3.1. *Retrospective and At-a-time Rationality*

There is little to say about *Retrospective Rationality* other than that it is a mistaken principle of rationality. While it may be appealing to think that a rational agent is one who is coherent and whose choices are optimal in a retrospective sense, this notion of rationality is useless to an agent who is wondering what to do. Such an agent will be making a decision at a time, and should be concerned to make the best decision possible *at that time*.

What matters when assessing a decision theory is just whether it recommends reasonable choices at any particular point in time, when an agent confronts the world and is wondering what to do. That is, only rationality at a time has any import when assessing

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

an agent's decision theory. So we regard Elga's discussion of classes of decision rules like SEQUENCE to be misguided, since he takes *Retrospective Rationality* to be important in assessing a decision theory.

This is not to say that the past and the future do not matter to decision-making at any particular point in time. Indeed, this is our second disagreement with Elga: he does not offer an adequate account of what *At-a-time Rationality* amounts to. The explicit points Elga makes on this issue are very limited: what an agent chooses now should not be affected by whether they were offered and refused a bet in the past or whether they were never offered that bet. Presumably, Elga means to condemn any decision rule that avoids the *NN* sequence in his sequential decision scenario by requiring that bet *B* be chosen whenever bet *A* is refused (even when it would otherwise be admissible to refuse *B*). On this particular point we agree with Elga. The agent's attitude to bet *B* should be the same in these two scenarios since in either case, choosing bet *B* is effectively ending up with the same outcome: *NB*.

Recall Elga's argument against SEQUENCE. In effect, what Elga argues is that rationality must impose the same constraints on you in each of the following situations:

- You have been offered and refused Bet *A*. You are now deciding whether to take Bet B.
- You have not been offered Bet *A*. You are now deciding whether to take Bet B.

We agree with Elga on this point: in both cases, choosing bet *B* means effectively ending up with *NB*. But Elga is, at least implicitly, committed to something much stronger as regards *At-a-time Rationality*. Elga's analysis suggests that past choices and future choices are always irrelevant to an agent's current choice amongst the options available to them. This implies that your current decision about whether to accept bet *B* should be unaffected by whether or not you previously accepted bet *A*. To see how this plays out, consider the following two situations:

- You have been offered and refused Bet *A*. You are now deciding whether to take Bet B.
- You have been offered and accepted Bet *A*. You are now deciding whether to take Bet B.

It does not follow from what Elga has argued above that you should choose the same way in these situations, but Elga does seem to be committed to the rational constraints on choice being the same in these two cases. Note that choosing bet *B* in these two cases means effectively choosing different packages of bets: *NB* in the first as opposed to *AB* in the second.

Now consider the case of future choices.

- You are deciding whether to take bet A in the knowledge that Bet B will be offered right after.
- You are deciding whether to take bet A in the knowledge that no future bets will be offered.

Should rationality require the same things of you in this case? Again it does not follow from Elga's argument that this should be so and yet Elga seems committed to its being so. Again, from his limited argument that past choices should sometimes be irrelevant

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

to decision, Elga apparently assumes that past choices and predicted future choices should always be irrelevant to decision.

In Section 4 our own position regarding *At-a-time Rationality* will be made clear. At this point, let us just say that we hold that choices *should* be affected by past and future decisions. This is because both past choices and (beliefs about) future choices impact on what the agent expects their current choice to deliver in terms of the final outcome.

3.2. *The strong package principle*

Taken together, Elga's commitment to *Retrospective Rationality*, his apparent interpretation of *At-a-time Rationality* and his assumption that the two must march in lockstep, yield a position we name the Strong Package Principle. We have criticised the assumptions that entail this principle, but let us now look at the principle itself.

Note first that the Strong Package Principle requires that admissibility for bets in isolation should agglomerate straightforwardly. This means that even the relation of 'not dispreferring' should agglomerate straightforwardly. (In the imprecise context, a bet is admissible if it is not dispreferred to any other available option, so if admissibility must agglomerate straightforwardly, so too must not dispreferring.) For Elga's decision problem, if the agent does not disprefer N to A and does not disprefer N to B, then the strong package principle requires that the agent does not disprefer NN to AB.[5] But the latter claim is not reasonable since NN is dominated by AB. So the Strong Package Principle effectively requires that one or both of the former two claims does not hold, that is, N must be dispreferred to at least one of A or B. This greatly constrains the class of acceptable decision rules for handling imprecision.

It is worth noting that, as our label suggests, the Strong Package Principle is reminiscent of a similar condition, which we call the 'Weak' Package Principle, underlying the synchronic Dutch book argument.[6] This is the principle that bets deemed weakly preferable in isolation must be weakly preferred as a package. Some have argued against this principle (see especially Schick 1986), claiming that it is an artificial constraint on actual choice situations; the Dutch book argument does not show that an ordinary agent with incoherent beliefs would necessarily suffer a sure loss, even if they were not averse to betting and it was inevitable they would meet a crafty bookie. The savvy agent would rather assess packages of bets differently from bets in isolation. It is only under the artificial constraint that bets must be assessed in isolation that the agent is vulnerable to sure loss.

If the Weak Package Principle is open to criticism, then so much the worse for the stronger principle. The added logical strength of the latter principle comes by its

5 The Strong Package Principle is thus *not satisfied* by the imprecise choice rule we mentioned earlier, which holds that an option is admissible/not dispreferred if and only if it is not dominated by another available option. The relation of domination does not agglomerate straightforwardly: neither A nor N dominates the other, likewise for B and N, but the sequence NN is dominated by AB.

6 Of course, the imprecise probabilist must reject something of the synchronic Dutch book argument for (precise) probabilism. But note that the Weak Package Principle need not be rejected. Imprecise probabilists have offered a variety of ways out of the Dutch book theorem. Paris (2005 [2001]) offers a Dutch book argument where the underlying logic is non-classical: upper and lower probability models come out as the way to avoid sure loss. Bradley (2012) argues that if the agent is not willing to take either side of some bets, then the agent need not be a precise probabilist. Both these arguments endorse the Weak Package Principle.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

constraining the relation of admissibility or 'not dispreferring', which is stronger than merely constraining the relation of preference.

In the precise context, agents are required to have complete preferences, so 'not dispreferring' amounts to the same thing as 'weak preference'. It is only in the imprecise context that the two come apart. This goes some way to explaining how this subtlety – that Elga's package principle is stronger than the 'original' principle in the Dutch book argument – can get overlooked.

4. ANALYSING SERIES OF BETS: SOPHISTICATED CHOICE

In this section we give our own analysis of Elga's betting problem, which amounts to taking the *sophisticated choice* approach to sequential decision making.[7] The starting point is an appropriate representation of Elga's problem, as per Figure 2. It allows us to highlight some important points about *At-a-time Rationality* in the context of sequences of decisions. To begin with, notice that the final outcomes (the four rightmost nodes of the figure) are *packages of bets* that result from the sequences of individual choices. We will see that the past and future choices matter to present choice insofar as they affect the overall outcome that you expect to get in making your present choice.

Let us elaborate the sophisticated-choice approach. It is clear from Figure 1 that the first choice is not simply a choice between Bet *A* and the status quo. The first choice is a choice between one of two future choice scenarios. You are contemplating the choice between Bet *A* and the status quo in the knowledge that you will be offered the Bet *B*/status quo choice soon after. We can reinterpret your choice at the initial node as a choice between different future choices; as a choice between:

- A choice between *AB* and *AN* (Up).
- A choice between *NB* and *NN* (Down).

How should we assess these future choices? What is relevant to a choice between choices? How do you evaluate an option that is itself a choice? Standard backwards induction holds that a choice can be evaluated in terms of its best option. This is because that's what you'd choose, if you chose that choice. Consider, for instance, the following decision problem: Four cards are placed in two envelopes. Envelope 1 contains a card that says '£10' and a card that says '£1'. Envelope 2 contains a card that says '£8' and a card that says '£9'. You are offered the choice between the two envelopes on the understanding that you will then get to choose which of the cards you get. You win the amount of money on the card you choose.

1. Envelope

   (a) £10
   (b) £1

7 Moreover, we consider sophisticated choice to be the orthodoxy when it comes to analysing sequential-decision problems.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

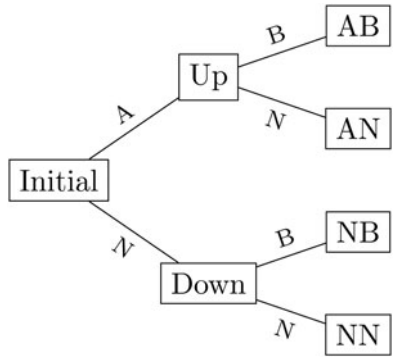


Fig. 2. Elga's problem as a decision tree

2. Envelope 2

(a) £8
(b) £9

Should you choose 1 or 2? You should choose 1, because 1a is better than any choice you could make if you chose 2. You work backwards, considering what you would do in each possible situation. At the initial node, the value of a choice is the value of what you would end up choosing. So at the initial choice between envelopes, you treat envelope 1 as being as good as you consider the future choice resulting from choosing that envelope to be.

So does this help to solve the Elga problem? Not obviously, because the two options in both sub-trees are incommensurable. Refer to Figure 2. At the initial node, choosing 'up' leads to a choice between *AB* and *AN*, and neither has maximal value: they are incommensurable. They have expected values of {5} and [−10, 15] respectively. Likewise for initially selecting 'down' where *NB* and *NN* have values [−10, 15] and {0} respectively. So when evaluating the 'up' option, it is not clear what its value should be, since it is still not clear what you would do, were you to end up in this situation.

In what follows, we consider two choice rules for decision-making in the imprecise context. Both rules are generalisations of standard expected utility theory,[8] and although the rules do not exhaust the proposals for decision-making under imprecision, they will be sufficient for illustrating the claims in this paper. The two rules are *gamma-minimax*, which says that admissible options are those available options that have maximal worst-case expected utility, and a *non-dominated-set* (*NDS*) rule, which says that admissible options are those available options that are not *weakly dominated, in terms of expected utility*, by any other options.[9]

8 That is, in the case that there is no imprecision and the agent in fact has precise subjective probability and utility functions, the rules equate to standard expected utility theory.

9 More precisely, an option A_i is *weakly expected utility (EU)-dominated* by another option A_k if, for all probability distributions in the agent's representor, $EU(A_k) \geq EU(A_i)$, and for at least one probability distribution in the agent's representor $EU(A_k) > EU(A_i)$. The proposed NDS rule is similar to the so-called *Sen–Walley maximality* rule; the difference is that dominance for the Sen–Walley rule is strict dominance, whereas for our NDS rule it is weak dominance.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

Consider the gamma-minimax rule first. You work backwards from your future choices. At the 'up' node, *AB* would be chosen because it does better on the gamma-minimax criterion. Gamma-minimax amounts to valuing acts by the lower end of their associated intervals of expectations. So *AB* is valued at 5, *AN* at −10. At 'up', you would opt for *AB*. At 'down', *NN* would be chosen since *NN* is valued at 0 to *NB*'s −10. Now each option is evaluated with respect to your current values – that is, your initial values – in terms of the outcome your future choices would lead to. At the initial node, then, going 'up' is effectively valued at the expected utility of *AB*, since that's what you will choose if you arrive at the 'up' node. And going 'down' is valued at what you would choose there: *NN*. Gamma-minimax values a future choice of *AB* at 5 compared with 0 for *NN*. So in fact *NN* would never be chosen by gamma-minimax. This differs from Elga's analysis, since he evaluates the first choice in isolation, and not as a choice between future choices.

This is not an entirely satisfactory resolution of the problem, however, because gamma-minimax is problematic (see, for example, Seidenfeld 2004). Indeed, we will expose a shortcoming of this rule at the end of the paper. But for now, note that the gamma-minimax sophisticated chooser would not end up with the package *NN*. This already casts doubt on Elga's claim that no plausible imprecise decision rule makes *NN* impermissible. This point has been made by Chandler (in press) and Sahlin and Weirich (2014).[10]

The more widely accepted of our pair of rules is the non-dominated-set (NDS) rule which we appealed to earlier. This rule is more in the spirit of imprecision because it does not 'induce a preference' in cases where options are simply incommensurable. At the later choice nodes, whether you choose 'up' or 'down', both options are admissible. That is, in neither case – up or down – do either of the options dominate the other. It is not straightforward or widely agreed how to implement sophisticated choice in this kind of situation. That is, how should these indeterminate future choices get evaluated from the position of the initial node? You are sat at the initial node, speculating about your future choices. You think 'If I were to go 'up', I would have a choice between *AB* and *AN*, and I will find those options incommensurable'. This doesn't seem a particularly satisfying resolution of your future choice. How ought you compare such an indeterminate choice with another indeterminate choice that would result from choosing 'down'?

This is our suggestion: the expected utility profiles of both 'up' and 'down' are calculated by considering the expectation of these options for each probability distribution in the agent's representor, in the knowledge that either of the two available options may be chosen at the later choice nodes. Colloquially, this is how you would think through your position at the initial node: for each sharp probability in your set of probabilities you think 'If my probability were sharp and given by Pr, I would value the future options by their (precise) expectations with respect to Pr. I know that my future choice is actually indeterminate – because my belief is imprecise – so I must evaluate the option of going 'up' by how I value all of the possible resolutions of my indeterminate future choice.'

Take, for instance, one of the extreme probability distributions in your representor: $\Pr_1(H) = 1$. For this distribution, 'up' has an expected utility profile of $\{5, -10\}$, since either

10 Thanks to Wlodek Rabinowicz for pointing us to this latter paper. Elga has acknowledged that he was mistaken on this point in an erratum to his paper that is available from his website.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

AB or *AN* may eventuate and Pr_1 evaluates those outcomes at 5 and −10 respectively, and 'down' has an expected utility profile of {15,0}, since here either *NB* or *NN* may eventuate. Every probability function in your set of probabilities yields such a set of utilities that represents how that probability evaluates the options of 'up' and 'down'. At the initial node, then, choosing 'up' is evaluated in terms of a set of sets of expected values, one for each function in your representor. For every $p \in [0,1]$ there is a set of expectations $\{5, -10p + (1-p)15\}$. This reflects one probability function's view on the value of the indeterminate choice that 'up' amounts to. *NN* is still a permissible option on this understanding since some probability functions in your representor give *NN* higher expected utility than *NB*, thus *NN* is not dominated by something in the branch of the decision tree where it is an option.

The question that arises is whether an imprecise decision rule is irrational if the agent in question anticipates, for Elga's decision problem, that *NN* may be chosen. This is subtly different from the issue of *Retrospective Rationality*, because it is rather about *anticipating* the possibility of ending up with a package of bets that is dominated by another package of bets that would result from a different set of choices.

Some authors flatly state that it is not irrational for a sophisticated agent to anticipate making a set of choices that is dominated by another set of choices in the decision tree (see especially Seidenfeld 1994; also Rabinowicz 1995). The argument for this position is that you never choose an option that is dominated relative to those that are actually available to you at the time. To use Seidenfeld's terminology, some strategies in the decision tree are simply not available; they are not *dynamically feasible*. You must first consider what you will choose at future times, and at these times, it may be that a dominated option is chosen, like *NN*, because the dominating option is not then part of the choice set. Working backwards: at the earlier, initial choice point, you must acknowledge what your future self will choose. So it may come about that you pursue a strategy that is dominated by another strategy that is unfortunately not possible due to inevitable future choices.

We are less permissive on the issue of anticipating that you will make a sequence of choices that is dominated by another sequence of choices in the decision tree. We accept that the *mere possibility* of ending up with a dominated strategy relative to the entire tree is *not* irrational but we hold that knowing you will end up with a dominated strategy is irrational. To take Elga's example again, if you know that you will end up with *NN*, then your decision theory is defective. But if your decision theory is such that *NN* just *might* come about, but then again might not because other strategies are also admissible at the relevant choice points, then you do not anticipate a sure loss. And typically the sign of irrationality is a *sure* loss and not a *possible* loss.[11] The upshot here is that the NDS decision rule is not irrational in virtue of permitting but not mandating pursuing the sequence of choices *NN*.

11 Moreover, as others have argued in the context of *non-forcing* money-pumps (see Gustafsson 2010), what is rationally permissible may not ever be chosen if further extra-rational principles are brought to bear on a decision. For instance, the agent's secondary principle might stipulate that amongst admissible options, an option that is dominated by another albeit unfeasible strategy in the decision tree should not itself be chosen.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

5. A BETTER DUTCH BOOK VARIANT

Elga positions his argument against imprecise probabilism as a 'variant of the diachronic Dutch book argument'. Here we have argued that his assumptions are too strong to warrant that label. The discussion of the previous section, however, suggests an alternative and better variant of the diachronic Dutch book argument, which will adjudicate against *some*, but not *all* decision theories for handling imprecision. That is, we endorse Elga's basic approach: your decision theory – your belief state and decision rule – should answer to questions of practical rationality. Making obviously bad sequences of decisions speaks against a decision theory. We have claimed that Elga's constraints on sequential decision making are problematic. Our own analysis of Elga's betting scenario gestures towards a better criterion for adjudicating among imprecise decision theories: a better variant of the Dutch book argument.

In short, the better Dutch book variant is this: when preferences and beliefs are assumed to remain stable, it is irrational to be *sure* of pursuing a strategy which is dominated by another strategy in the sequential decision tree. The possibility of ending up with a dominated strategy relative to others in the tree is, however, rationally permissible (cf. Steele 2010: 471–4). This Dutch book variant adjudicates against one of the two choice rules outlined in this paper – the gamma-minimax rule. While gamma-minimax looks to 'get it right' for Elga's problem in Figure 2, in other decision scenarios, it recommends a sure loss.[12] The NDS rule, on the other hand, leads only to possible losses,[13] and so is satisfactory by the standards of our Dutch book variant. Moreover, since there is no obligatory sure loss, such rules can be combined with extra-rational principles of choice, so that an agent may never choose an (admissible) dominated option. These sorts of principles are currently under-explored, in our opinion, and are deserving of further consideration.

We have argued that Elga's constraints on rational imprecise choice are too strong. Here is another way to see that Elga's argument must be too strong. Note that the only feature of imprecise probabilities that really makes an appearance in Elga's argument is that they give rise to incommensurable sets of expectations. This, in itself, is not a criticism: all pragmatic arguments ultimately rely on facts about the expected utilities and choices of agents. However, it does not matter to Elga's argument that the source of the incommensurability is imprecise belief: it could just as well have come from imprecise utilities. So Elga's argument is not so much an argument against imprecise probabilities, but rather an argument against any kind of incommensurability in sequential choice. That is, any decision problem that involves incommensurability of options is equally susceptible to Elga's argument.

Consider a case where there are two goods that are objectively incomparable, or goods whose utilities are objectively incomparable. Perhaps Gross Domestic Product and biodiversity are two such goods.[14] Consider a choice between a certain increase in GDP X and a certain increase in biodiversity Y. You have no preference between them. Now consider two further goods, $X+\varepsilon$ and $Y+\varepsilon$ which are such that they are strictly preferred to X and Y respectively, but incomparable to each other and to the other goods. So $X+\varepsilon$ is a

12 For example, Steele (2010) offers an example decision problem on p. 473. The example is set up for a slightly different decision rule, but it works in this case too.

13 This is just a conjecture at present; we do not offer a proof here.

14 We are not arguing for the plausibility of such incommensurabilities: we merely want to point out that Elga's argument rules out these things.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press

slightly higher level of GDP, $Y+\varepsilon$ is a slightly more diverse biosphere. Now, you get to choose whether I add the ε to the Y or the X. You then get to choose between the goods. So you choose whether you have a choice between X and $Y+\varepsilon$ or between $X+\varepsilon$ and Y. Whichever you choose, you end up with a choice between two incommensurable options. That is, you have no preference between X and $Y+\varepsilon$. So it is permissible to choose X despite the fact that it is dominated by something on the other branch of the tree (namely $X+\varepsilon$). Thus an argument exactly analogous to Elga's seems to lead to the conclusion that goods should be commensurable. This suggests that Elga's argument is stronger than he makes it out to be. This need not trouble those who think that all goods are commensurable, but those who think that incommensurability of goods is at least a possibility should perhaps be sceptical of Elga's line of thought. In short, Elga's argument does not just rule out imprecise belief, but also any kind of imprecise value. Since imprecise values are arguably more plausible than imprecise beliefs, this tells against Elga's argument: he is arguing against something that is actually quite a plausible component of decision theory.

To conclude, if we allow our imprecise probabilists sophisticated choice, and do not judge them by unreasonably strong standards of rationality, they need not make unavoidably bad sequences of decisions. Thus imprecise probabilistic beliefs are not irrational.

REFERENCES

Bradley, S. 2012. 'Dutch book Arguments and Imprecise Probabilities.' In D. Dieks, W. J. González, S. Hartmann, M. Stöltzner and M. Weber (eds), *Probabilities, Laws and Structures*, pp. 3–17. New York, NY: Springer.

Chandler, J. In press. 'Subjective Probabilities need not be Sharp.' *Erkenntnis*.

Elga, A. 2010. 'Subjective Probabilities should be Sharp.' *Philosophers' Imprint* 10.

Gustafsson, J. E. 2010. 'A Money-pump for Acyclic Intransitive Preferences.' *Dialectica*, 64: 251–7.

Paris, J. 2005 [2001]. 'A Note on the Dutch book Method.' In *Proceedings of the Second International Symposium on Imprecise Probabilities and their Applications*, pp. 301–306.

Rabinowicz, W. 1995. 'To Have One's Cake and Eat it too: Sequential Choice and Expected-utility Violations.' *Journal of Philosophy*, 92: 586–620.

Sahlin, N.-E. and Weirich, P. 2014. 'Unsharp Sharpness.' *Theoria*, 80: 100–3.

Schick, F. 1986. 'Dutch Bookies and Money Pumps.' *Journal of Philosophy*, 83: 112–19.

Seidenfeld, T. 1994. 'When Normal and Extensive Form Decisions Differ.' *Logic, Methodology and Philosophy of Science*, IX: 451–63.

——. 2004. 'A Contrast between two Decision Rules for use with (Convex) Sets of Probabilities: Γ-maximin versus *E*-admissibility.' *Synthese*, 140: 69–88.

Steele, K. 2010. 'What are the Minimal Requirements of Rational Choice? Arguments from the Sequential Setting.' *Theory and Decision*, 68: 463–87.

van Fraassen, B. 1990. 'Figures in a Probability Landscape.' In M. Dunn and K. Segerberg (eds), *Truth or Consequences*, pp. 345–56. Amsterdam: Kluwer.

SEAMUS BRADLEY is a Postdoctoral Research Fellow at the Munich Centre for Mathematical Philosophy, LMU. His main interests are formal epistemology and decision theory, especially imprecise probabilities; and philosophy of science.

KATIE STEELE is an Associate Professor in the Department of Philosophy, Logic and Scientific Method at the London School of Economics and Political Science. Her work chiefly concerns rational-choice theory, methodology of science, and evidence-based policy.

https://doi.org/10.1017/epi.2014.8 Published online by Cambridge University Press