Noûs

NOÛS 00:0 (2016) 1–19
doi: 10.1111/nous.12148

Imprecise Probabilities and Unstable Betting Behaviour[1]

ANNA MAHTANI
The London School of Economics

Abstract

Many have argued that a rational agent's attitude towards a proposition may be better represented by a probability range than by a single number. I show that in such cases an agent will have unstable betting behaviour, and so will behave in an unpredictable way. I use this point to argue against a range of responses to the 'two bets' argument for sharp probabilities.

Section 1: Introduction

Are subjective probabilities always precise? Plausibly we have *some* precise probabilities: to take the sort of clean case you might find in a probability textbook, your probability that a fair coin will land heads on its next toss is presumably precisely 0.5. But once we turn our attention to a wider class of claims, then the idea that our subjective probabilities are *always* precise seems implausible. What is your probability, for example, that (R) it will rain at midday in London on New Year's Day 2025? Perhaps you are neither very confident that it will rain, nor very confident that it won't, in which case your probability should be somewhere around 0.5. But what *exactly* is your probability in R? Is it, for example, 0.424? Or is it very slightly higher—say 0.425? You might be able to state a number if pressed—but the number you state will be arbitrary: you won't know that this is your probability in R.[2] Furthermore, it is not clear what could possibly make it the case that your probability is one particular number (say, 0.424) rather than another (say, 0.425). Many theorists have responded to this sort of case by claiming that subjective probabilities—even the subjective probabilities of rational agents—need not be precise. On one prominent version of this thought (defended by Levi 1974, Jeffrey 1983, van Fraassen 1990, Bradley 2009 and Joyce 2011, and amongst others) a person's doxastic state is better represented by a *set* of probability functions, rather than by one unique probability function. Your probability in a particular claim can then be thought of as a range (the range of numbers that the probability functions in your set assign to that claim) rather than a single number. Thus subjective probabilities need not be sharp.

The claim that rational agents may have imprecise subjective probabilities has faced a variety of objections. In this paper I focus on an objection concerning sequential choice. The problems that advocates of imprecise probabilities face in cases of sequential choice have been discussed in the literature from both economics and philosophy.[3] Here I will focus on one particular version of this objection—the 'two-bets argument', recently championed by Adam Elga (Elga 2010)—though many of the points that I make will also apply to other sequence based arguments for sharpness.[4] Elga's argument has attracted plenty of interest, and there has been a recent frenzy of papers pushing a particular sort of objection: the key thought is that Elga has overlooked both a range of sophisticated decision rules and the resources of game theory, and that once these are taken into account the argument does not go through. This point has been made in different forms by Teddy Seidenfeld, Nils-Eric Sahlin and Paul Weirich (Sahlin & Weirich 2012), Jake Chandler (Chandler 2014), Katie Steele and Seamus Bradley (Bradley & Steele 2014).[5] Elga in his errata (Elga 2012) addresses Seidenfeld's version of this objection, and his come-back is to gesture towards some of the short-falls of the particular decision rule that Seidenfeld has in mind. However, this response will only work against that particular version of the objection: the objection can be made in a similar form using a wide range of different decision rules—a range which is still proliferating. To respond to every version of the objection in a similar way, a champion of the two-bets argument would need to counteract each decision rule as it is put forward. In this paper I offer a simple powerful response that should allow us to dismiss all versions of this objection to the two-bets argument at a stroke. I argue that an agent with an imprecise subjective probability will display unstable betting behaviour, and this fact is at odds with the common strategy of these objectors. You can see my work in this paper as a giving a strengthened version of the two-bets argument which is immune to a range of objections. I believe that there may yet be a response even to this strengthened version (and so a defense of imprecise subjective probabilities)—but this does not come cheap.

I begin in the next section by setting out Elga's version of the two-bets argument against imprecise subjective probabilities.

Section 2: The Two-Bets Argument

Suppose that you have an imprecise probability in claim R. We can suppose that your probability in R is the range [0.1–0.8]: the exact numbers do not matter here, as the discussion below can be adjusted to apply to any range you like. Now for the two-bets argument. You are offered the following 'great series of bets' (Elga 2010, p. 4)—great because taken together they guarantee you a profit:

> Bet A: If R is true, you lose \$10; otherwise, you win \$15.
> Bet B: If R is true, you win \$15; otherwise, you lose \$10.

You know that you will be offered both bets in succession, and that nothing significant will happen between your being offered bet A and your being offered bet B—i.e. you will not gain any evidence, or have any great insight, or change in any other relevant way. You value money linearly (assume that \$1 is one utile), and let's assume that there is nothing else that you value of any relevance for this scenario. Elga argues that an agent who rejected both bets would be irrational. After all, why reject both bets, ensuring that you win nothing, when you could accept both bets and make a certain \$5? I agree with Elga that to reject both bets would be irrational—and notice that no argument is required here, nor any special technical sense of 'irrational'. Rejecting both bets would just be a stupid thing to do—not something that *you* would do, for example.[6]

We can grant then that rejecting both bets would be irrational. But which of the two bets are you rationally required to accept? If you had a precise probability (x) in R, then there would be an obvious way to answer this question. You ought, rationally, to carry out actions which maximize expected utility.[7] Should you then accept or reject bet A? The expected utility of accepting bet A is $(-10x + 15(1\text{-}x)) = 15 - 25x$, and the expected utility of rejecting bet A is 0. Thus if $x < 3/5$, then you should accept A; if $x > 3/5$, then you should reject A; and if $x = 3/5$, then you are rationally permitted to either accept or reject A. Should you accept or reject bet B? The expected utility of accepting bet B is $(15x - 10(1\text{-}x)) = 25x - 10$, and the expected utility of rejecting bet B is 0. Thus if $x > 2/5$, then you should accept bet B; if $x < 2/5$, then you should reject bet B; and if $x = 2/5$, then you are rationally permitted to either accept or reject B. We can see then that whatever your probability in R, you are rationally required to accept at least one bet: if $x < 3/5$, then you are rationally required to accept bet A; and if $x > 2/5$, then you are rationally required to accept bet B; and of course x (i.e. your probability in R) must be either less than 3/5, or greater than 2/5 (or both).

But what should you do if you have an imprecise probability in R? In particular, what should you do if your probability in R is the range [0.1–0.8]? What do the rules of rationality dictate here? Where probabilities are sharp, we have the rule 'maximize expected utility', which constrains the actions that can rationally be taken by the agent. What is the equivalent where probabilities are imprecise? As Elga writes, 'Anyone who claims that [rational] agents can have unsharp probability functions owes a corresponding account of how unsharp probabilities constrain rational action' (Elga 2010, p. 3). Elga then explores a number of possible ways of meeting this challenge, and finds that each is inadequate. Throughout, this compelling claim is taken for granted: it would be irrational to reject both bets. Thus any attempt to state the 'rules that constrain rational action' will be judged incomplete unless they entail that it is irrational to reject both bets.

We can immediately eliminate any 'permissive choice' rules according to which it is rationally permitted to either accept or reject the first bet, and then (regardless of your choice over bet A) rationally permitted to either accept or reject the second bet. The underlying idea behind these permissive choice rules is this: an action is rationally permitted provided that it maximises expected utility relative to some probability function in the set that represents your doxastic state. On this view, you are rationally permitted to accept bet A, because for some probability functions in your range—such as that which assigns 0.1 to claim R—the expected utility of accepting A $(((-10)(0.1)+(15)(0.9)) = 12.5$, where R is assigned 0.1) is higher than the expected utility of rejecting A (0 for all probability functions). You are also rationally permitted to reject A, because for some probabilities in your range—such as that which assigns 0.8 to R—the expected utility of accepting A $(((-10)(0.8)+(15)(0.2)) = -5$, where R is assigned 0.8) is lower than the expected utility of rejecting A. Similar reasoning shows that you are also rationally permitted to either accept or reject bet B. Thus on this view you are rationally permitted to accept or reject either bet—and in particular you are rationally permitted to reject both bets. This violates the compelling claim that no rational agent would reject both of these bets, and thus it seems we can eliminate all 'permissive choice rules'.

Good objections can be raised to a host of other candidate rules for rationality. Elga rejects the 'mid-point rule'—that if rational you will act as though you had a sharp probability at the mid-point of your probability range—on the grounds that this 'robs unsharp probabilities of their point' (Elga 2010, p. 6). Elga rejects the 'narrowing proposal'—according to which your probability range in a claim should narrow in line with the actions that you take—on the grounds that it cannot be a rational requirement that your probability should change without a change in your evidence. He rejects the 'planning proposal'—according to which in the two-bets scenario a rational agent would make a plan and stick to it—on the grounds that it cannot be a brute rule of rationality that agents should stick to their plans.[8] And finally Elga rejects the 'sequence proposal'[9]—according to which an agent's choices throughout a *sequence* of bets must maximise expected utility relative to some probability function in the agent's range—on the grounds that this rule would have the consequence that an agent offered merely bet B may rationally refuse it, but an agent in the two-bet scenario who has refused bet A may not refuse bet B, even though the choice that the agent faces in these two situations are the same in every relevant respect.[10] It seems then that the advocate of imprecise subjective probabilities has no good answer to the question 'how do unsharp probabilities constrain rational action?'. Without a good answer to this question, the defender of imprecise probabilities is in an untenable position, and so it seems we must conclude that rational agents have only sharp probabilities.

This argument has faced a barrage of criticism—including the prevalent objection that I focus on in this paper, raised by Seidenfeld, Sahlin & Weirich (2012), Bradley and Steele (2014) and Chandler (2014). These authors point to the range of sophisticated decision rules developed by imprecise probability theorists, and the resources of game theory: once these are taken into account (these authors claim), the two-bets argument no longer goes through. In the next section I run through a version of this point, before turning to my response in section 4.

Section 3: The Objection to the Two-Bets Argument

Imprecise theorists have developed various decision rules, including maximin (Gärdenfors & Sahlin 1982 and others), maximax (Satia & Lave 1973), Ellsberg's rule (Ellsberg 1961) and Hurwicz's criterion (Arrow & Hurwicz 1972). Elga does not focus on these rules in his 2010 paper,[11] so let us consider how these rules fare in his two-bets scenario, taking maximin as our focus.

The rule maximin works as follows. An agent's doxastic state is represented by a set of precise probability functions, and when considering a possible action, there is an expected utility for that action relative to each precise probability function in the agent's set. Amongst these expected utilities for the action, one will be the lowest—and so each action has a *minimum* expected utility. According to maximin, when faced with a choice, a rational agent will carry out whichever action has the maximum minimum expected utility.

To see maximin in action, suppose first that you are presented with just bet A: if R is true, you lose \$10; otherwise, you win \$15. Suppose as before that your subjective probability in R is the range [0.1–0.8]: in other words, your probability function is a set of precise probabilities which between them assign all numbers in the range [0.1–0.8] to claim R. Accepting bet A has an expected utility relative to each of these precise probabilities. For example, the probability that assigns 0.1 to claim R, gives the act of accepting bet A an expected utility of $((0.1)(-10) + (0.9)(15)) = 12.5$; the probability that assigns 0.5 to claim R gives the act of accepting bet A an expected utility of $((0.5)(-10) + (0.5)(15)) = 2.5$. The probability that gives the lowest expected utility to the act of accepting bet A is the probability that assigns 0.8 to claim R: this gives the act of accepting bet A an expected utility of $((0.8)(-10) + (0.2)(15)) = -5$. Thus the minimum expected utility of accepting bet A is -5. The expected utility of rejecting bet A is the same for each probability function: it is simply 0. Thus the minimum expected utility of accepting bet A (i.e. -5) is lower than the minimum expected utility of rejecting bet A (i.e. 0), and so maximin rules that if you are rational you will reject bet A—for this is the act with the maximum minimum expected utility. We can use similar reasoning to show that the maximin rule requires you to reject bet B if offered in isolation: the expected utility of accepting bet B will be at its lowest when we take the probability of R to be 0.1, and so the minimum expected utility of accepting the bet is $((0.1)(15) + (0.9)(-10)) = -7.5$, and this is less than the minimum expected utility of rejecting bet B (i.e. 0). It may be tempting then to assume that the maximin rule requires you to reject both bets in the two-bets scenario—and so that the maximin rule conflicts with the plausible claim that an agent who rejects both bets is irrational.

But in fact, as Bradley & Steele (2014), Chandler (2014), Sahlin & Weirich (2012) and Seidenfeld have all pointed out—and as Elga has conceded (Elga 2012, p. 1)—an agent who is obeying the maximin rule will not reject both bets.[12] To see why this is, we need to think about the two-bets scenario like a game theorist. In the two-bets scenario you know before you are offered bet A that you will also be offered bet B. You also know that nothing significant will have changed between your being offered bet A and your being offered bet B. Thus if your behaviour is in line with maximin when you are offered bet A, then it will also be in line with maximin when you are offered bet B. If we suppose that you know from the start that your behaviour will be in line with maximin throughout, and you also know your preferences and doxastic state, then at the time when you are offered bet A you can figure out how you will choose when faced with bet B. Whether you accept or reject bet B will depend on the minimum expected utilities of these two actions, which are laid out in the table below:

		Minimum Expected Utility
If A has already been accepted	Accept B	5
	Reject B (leaving yourself exposed to just bet A)	$((0.8)(-10)+(0.2)(15)) = -8+3 = -5$
If A has already been rejected	Accept B	$((0.1)(15)+(0.9)(-10)) = 1.5-9 = -7.5$
	Reject B	0

You can see then that if bet A is accepted, then bet B will also be accepted (because $5 > -5$), and if bet A is rejected then bet B will also be rejected (because $0 > -7.5$). You can figure all of this out before you are offered bet A, and so when deciding whether to accept bet A, you do so with the knowledge that if you accept bet A you will end up accepting both bets, and if you reject bet A you will end up rejecting both bets. Thus you are effectively deciding whether to accept both bets (with a minimum expected utility of 5), or to accept neither bet (with a minimum expected utility of 0). Given that you should carry out the act with the maximum minimum expected utility, you should accept bet A. Thus maximin rules that a rational agent (who has a probability range of [0.1–0.8] in R, and who values only money, and values it linearly) will end up accepting both bets. It seems then that the imprecise probability theorist *can* answer the challenge to produce a plausible rule for rational action that gives an acceptable result in the two-bets scenario.

Notice here that the maximin rule requires that you accept bet B once bet A has been offered and accepted, but that you reject bet B if is offered in isolation. Similarly maximin requires you to accept bet A if you know that bet B will also be offered, but to reject bet A if it is offered in isolation. It might seem then that maximin should face the sorts of objection that Elga levels at the sequence proposal. According to the sequence proposal, you are rationally permitted to reject bet B if offered in isolation, but not rationally permitted to reject bet B if you have already rejected bet A—for this would be an irrational *sequence* of actions. In objection to this proposal, Elga protests that rationality can't impose different requirements on you in these two cases, since in each case you '[face] choices that are exactly the same in every respect that [you care] about' (Elga 2010, p. 8). This seems like a strong objection to the sequence proposal: can a similar objection against the maximin rule carry the same force? Here we might compare two cases: in one case you are offered bet B in isolation; in the other case, you are in the two-bets scenario, have just accepted bet A, and are now being offered bet B. Maximin rules (assuming as usual that your probability in R is the range [0.1–0.8]) that B should be rejected in the first case and accepted in the second case. Do you in these two cases face choices that are exactly the same in every respect that you care about? No: you care about minimum expected utility, and the minimum expected utility of the actions available to you are different in the two cases. Intuitively, accepting B in the first case is risky (with a minimum expected utility of -7.5) whereas accepting B in the second case is safely profitable (with a minimum expected utility of 5). Similarly maximin rules that if you are offered bet A in isolation you should reject it, but if you are offered bet A knowing that bet B will also be offered, you should accept A. Again, these scenarios are different in some respect that you care about: accepting A in isolation is risky (with a minimum expected utility of −5), but accepting A knowing that you will shortly accept bet B is safe (with a minimum expected utility of 5). Thus the objection that Elga levels against the sequence proposal cannot be applied to maximin.

Here it might be objected that we have not given the permissive choice rules a fair trial. When considering maximin, we assumed that when faced with bet B, you would take into account the choice already made over bet A, and we treated the utilities as cumulative. Furthermore, we assumed that when faced with bet A, you would take into account what choice you will make over bet B. What happens if we give the permissive choice rules the same treatment? We can start by figuring out whether you would accept bet B taking into account your earlier choice over bet A:

		Expected utility, relative to the probability function that assigns x to claim R.
If A has already been accepted	Accept B	5
	Reject B	$((x)(-10)+(1\text{-}x)(15))$
If A has already been rejected	Accept B	$((x)(15)+(1\text{-}x)(-10))$
	Reject B	0

Your doxastic state is represented by a range of probability functions, which between them assign R every number in the range [0.1–0.8]. Relative to some of these probability functions (e.g. that which assigns 0.8 to claim R), the expected utility, once A has already been accepted, of accepting bet B (5) is higher than the expected utility of rejecting bet B $(((0.8)(-10) + (0.2)(15)) = -5)$. Relative to some other probability functions (e.g. that which assigns 0.1 to claim R) the expected utility of rejecting bet B $(((0.1)(-10) + (0.9)(15)) = 12.5)$ is higher. Thus (according to the permissive choice rules) you may rationally accept or reject bet B once bet A has already been accepted. We can show similarly that you may either accept or reject B once A has already been rejected. Thus when considering whether to accept or reject bet A, you will not know whether you will subsequently accept or reject B. At least, you cannot figure this out from knowing your own probabilities and preferences, knowing that you are rational, and knowing that rational agents adhere to the permissive choice rules.

Thus the prevalent response to the two-bets argument cannot be made using the permissive choice rules rather than maximin. The advantage of maximin for these purposes is that it *determines* how you will respond to bet B: thus if you know that you are rational, and that rational agents adhere to maximin, and if you know your own probabilities and preferences, then at the point where you are deciding whether to accept bet A, you will be able to predict whether or not you will accept bet B (dependent on your decision to accept or reject bet A). This means that when deciding whether to accept bet A, you are effectively deciding between different action-bundles: you are choosing between accepting both A and B (for you know that if you accept A then you will also accept B), and rejecting both A and B (for you know that if you reject A then you will also reject B). Maximin rules in favour of accepting both bets—and so maximin can handle the two-bets scenario. Maximin is not the only rule that can handle the scenario in this sort of way: versions of the same argument can be given using maximax, Hurwicz's rule and others. What the rules have in common is that they are *not* permissive: they settle how the agent (if rational) will act—given the agent's preferences, doxastic state, and in some cases another measure such as the agent's optimism quotient.[13]

Thus we seem to have a strategy for responding to the two-bets argument, and that strategy works with a wide range of rules of rationality. One way a champion of the two-bets argument could respond would be to shoot down these rules one by one, and this is indeed how Elga responds in his errata to the objection as raised by Seidenfeld—by pointing out a problem with maximin (Elga 2012, p. 1). But Elga's strategy is problematic: firstly, because any problems with the individual rules must be weighed against the problems with the alternative—i.e. the counterintuitive claim that all subjective probabilities are sharp; and secondly because we can see how new more complex rules broadly along the same lines may yet be developed—and so finding objections to each one may prove a never-ending task.

I propose instead to object to the strategy itself. An essential part of the strategy is the claim that your decisions are rationally determined (given your doxastic state and so on), and so that your decisions are *predictable*. This is what makes it possible for you to approach the two-bets scenario as a game theorist: you can figure out which actions you will go on to take, and so you are effectively comparing bundles of actions from the start. I argue in the next section that this ability to predict your own actions is typically missing in cases where your subjective probabilities are imprecise.

Section 4: Unstable Betting Behaviour

Let's start by thinking again about claim R (the claim that it will rain in London on New Year's Day 2025). This serves as an example of a claim in which—intuitively—you have an imprecise probability. If you don't think it serves this purpose, then just replace throughout with some claim that does: all parties can agree that there are some claims in which you *seem* to have an imprecise probability rather than a precise probability. Now let's think about your betting behaviour—which bets you would accept and reject—concerning R. Take bets of this form:

D: You pay $\$v$, and you get \$1 back iff R.

Perhaps for some v, it will be obvious for you whether to accept or reject the bet. So perhaps you would gladly accept the bet where $v = 0.01$, and would reject the bet out of hand for $v = 0.99$. But for some v you may be unsure what to do. Perhaps for $v = 0.45$, for example, you waver. If pushed to make a decision, you may be able to do so, but your decision will have an arbitrary flavor. If you decide to accept the bet, for example, then you could easily have decided instead to reject it: though you accept the bet in the actual world, in a close possible world where you have all the same evidence and the same preferences as in the actual world, you instead reject the bet.

I claim that this sort of 'instability' in an agent's betting behaviour is typical of cases where intuitively we class an agent as having an imprecise probability.[14] For contrast, consider a case where R is some claim in which you have a precise probability: say that R is the claim that a particular coin that you know to be fair will land heads on its next toss. In this case your probability in R is precisely 0.5. If you are offered a bet of the form D, then whenever v is less than 0.5 you will accept the bet—and you will do so reliably: you will accept the bet in any close possible world where you have the same evidence and the same preferences.[15] Similarly whenever v is greater than 0.5, you will reliably reject the bet. It is only when v is exactly 0.5 that your betting behaviour is unstable: you might accept the bet or you might reject it—but you will do so on a whim and there will be relevant close possible worlds where you make a different choice. Where your probability in R is precise, then, unstable betting behaviour over a bet like D narrows to a point; where your probability in R is imprecise, unstable betting behaviour over a bet like D extends to a range corresponding to the range of numbers assigned to claim R by the probabilities in the set that represents your doxastic attitude. Thus in a case where your subjective probability in R is the range [0.1–0.8], you will reliably accept bet D for any $v < 0.1$, and you will reliably reject bet D for any $v > 0.8$, but for v such that $0.1 \leq v \leq 0.8$, your betting behaviour over bet D will be unstable: that is, there will be close possible worlds where you have your actual evidence and preferences where you accept the bet, and close possible worlds where you have your actual evidence and preferences where you reject the bet.

As an analogy, compare this case: you are briefly shown a large jar, filled to the brim with sweets. You are asked a question of form E, and told that if you give the right answer then you will win the jar of sweets:

E: Are there at least v sweets in the jar?

Perhaps where $v = 1$, you will confidently answer 'yes', and for $v = 1000{,}000$ you will confidently answer 'no' But for some v (perhaps for $v = 150$) you waver. Given the strong incentive to be in with a chance of winning the sweets, you will manage to cough up an answer, but the answer you give will be arbitrary. You will decide by tossing a coin, perhaps, or consulting your 'feelings', or you will allow yourself to be swayed by some irrelevant feature of your environment. The coin toss, your feelings, and the feature of the environment could all easily have turned out otherwise—in which case you would have given a different answer. Your behaviour here is unstable: if you say 'yes' in the actual world, then there is some close possible world where you have all the same evidence and preferences, but say no.[16]

Where behaviour is unstable, it is hard to predict. In the sweet jar example, with $v = 150$, an onlooker would have a hard time predicting how you would answer the question: this holds even if the onlooker knows exactly what evidence you have (e.g. perhaps (s)he shares your evidence and knows it), and knows how much you would like to win the prize. You yourself don't have enough information to *predict* how you will respond: you might *resolve* to respond a particular way, but that is a different matter. Similarly, to return to bet D, if you are offered a bet of this form with v falling somewhere within your subjective probability range [0.1–0.8], an onlooker—even one who knew about your evidence and preferences—may be unable to predict how you would respond. For you will just decide on a whim: perhaps you will toss a coin, or be influenced by some apparently irrelevant feature of your environment, or consult your feelings. Given that your decision can be swayed by such insignificant things, without knowing all of these factors—the outcome of the coin toss, the details of your environment, your biochemistry, and exactly how these factors will determine your decision—neither you nor an outside observer could reliably predict your betting behavior.[17]

The problem is that the prevalent strategy for responding to the two-bets argument relies on the assumption that you *can* predict your own betting behaviour. For a key move in the strategy is to think about the problem like a game theorist: we assume that at the point when you are offered bet A, you will be able to think ahead and figure out whether (depending on whether you accept bet A) you will accept bet B. But this assumption is unrealistic. In the cases we are interested in—the cases where intuitively subjective probabilities are imprecise—betting behaviour is typically unstable and so unpredictable.

To see the problem here, let us focus on the question of whether you would accept bet B, having already rejected bet A. If you have rejected bet A, then you have a choice between accepting bet B (paying out \$10, and getting \$25 back (i.e. winning \$15) if R is true), and rejecting bet B (losing and gaining nothing). Given that you value only money and value it linearly, we would expect your betting behaviour over this bet to be like that over a bet of the form D with $v = 2/5$ (i.e. you pay out \$2/5, and get \$1 back iff R). As v here is within your probability range for R ([0.1–0.8]), we would expect your betting behaviour over this D-form bet to be unstable and therefore unpredictable. Thus we would expect your betting behaviour over bet B—once you have rejected A—to similarly be unstable. As unstable betting behaviour is unpredictable, it would follow that you are unable to predict from the start of the sequence how you will respond to bet B if you reject bet A. Similar reasoning would show that your betting behaviour over B is unstable if you accept (rather than reject) bet A. Thus the prevalent strategy cannot be applied: we cannot use game theory to handle the two-bets scenario in the way that has been proposed.

Here a natural objection is that your betting behaviour would *not* be unstable or unpredictable in the two-bets scenario. After all, any rational agent (and you are hardly going to be deliberately irrational!) would not reject both bets. I heartily agree—but this is the very mystery that we are grappling with. The dialectic here is complicated, so I will pause to spell it out. We began with the uncontroversial piece of data that there are lots of everyday cases where people—ordinary, generally intelligent people, people who are 'rational' at least in the everyday sense—don't *seem* to have sharp probabilities, but perhaps have imprecise probabilities instead. An example is the case that I began with: your probability that (R) it will rain at midday on London in 2025. I claim that a mark of these sorts of cases is that the agent in question would display unstable betting behaviour across a range of bets—e.g. across bets of the form D with a range of values for *v*. That this unstable betting behaviour is at least a mark of these cases seems undeniable: after all, what motivation would we have for saying that an agent seems to lack sharp probabilities, if his or her betting behaviour was perfectly stable? Then the challenge is to handle the two-bets scenario. In this scenario, a rational agent ('rational' in the everyday sense—an ordinary, generally intelligent person such as yourself) would *not* reject both bets. How can this be explained? The prevalent response is to claim that rational agents act in line with some rule such as maximin, and further that they *know* they are doing so, and so can predict their future behaviour. If this is correct, then (as the advocates of this view have shown) a rational agent would not reject both bets in the two-bets scenario. There is no need (the advocates say) to rely on the moves that Elga argues against in his 2010 paper, such as the midpoint rule, or the planning proposal, or the sequence proposal. The sophisticated decision rule (e.g. maximin) together with game theory already has all the resources needed to explain why a rational agent will not reject both bets. The problem for this approach, of course, is that it relies on a false assumption: that even in cases where intuitively a rational agent has imprecise probabilities, that agent's betting behaviour is in line with some rule, and so is stable and therefore predictable. The very phenomena that we began with—and that advocates of the prevalent response to the two-bets argument want to uphold—is mischaracterized by this approach. Thus the prevalent response to the two-bets argument does not work.

It may seem as though my argument here has far-reaching (perhaps implausibly far-reaching) consequences. Given that I claim that the mark of an agent with imprecise probabilities is unstable betting behaviour across a range of bets (e.g. across bets of the form D with a range of values for *v*), does it follow that we should reject maximin altogether—and also maximax, the Ellsberg Rule, the Hurwicz criterion, and countless other sophisticated decision rules? After all, these rules prescribe how an agent should act—given that agent's doxastic state, preferences, and perhaps his or her optimism quotient or similar. While these remain fixed across close possible worlds, presumably the agent's behaviour will be stable—at least if the agent is rational. So don't all of these rules misrepresent the phenomena that we are interested in? The theorists who develop these rules of course think that a rational agent can have imprecise probabilities (otherwise they would not be motivated to develop the rules), but if we take instability of betting behaviour as a mark of an agent's having imprecise probabilities, then we have a problem—because these rules seem to leave no room for a rational agent's displaying unstable betting behaviour. I think that this thought is well worth pursuing, but in this paper I do not attempt to develop my point in this direction. After all, there are plenty of ways that advocates of sophisticated decision rules for imprecise probabilities might respond: they might say that there is no assumption that an agent's doxastic state, or preferences, or level of optimism will be stable in typical cases of imprecise probabilities; or they might point out that the agent's doxastic state is vague, so often no particular actions are determinately prescribed; or they might claim that the rules were only supposed to *model* rational behaviour, and should not be taken too literally. The problem that concerns us in this paper is that on the prevalent response to the two-bets argument the rules *are* taken literally, and the doxastic state, preferences and level of optimism are assumed to be determinate and stable. It is only if we take this view that we get a response to the two-bets argument: it is crucial that the agent knows, at the time when (s)he is offered bet A, how she will act when offered bet B, and this is not guaranteed if the agent's doxastic state is supposed to be unstable, or vague, or if the relevant rule for rationality (e.g. maximin) is a mere model to which rational agents need not literally conform. Thus while there may be various ways of reconciling rules such as maximin with the fact that where probabilities are imprecise, betting behaviour is unstable, there is no way of reconciling the prevalent response to the two-bets argument with this fact.

I have argued, then, that the prevalent response to the two-bets challenge is unsuccessful. In the next section, I suggest a strengthened form of the two-bets argument, and consider whether any attempts to refute the original argument could work against this strengthened form.

Section 5: The Two-Bets Argument Strengthened

I have claimed that unstable betting behaviour is a mark of imprecise probabilities. To be more precise, my claim is that where an agent has a probability [*m*–*n*] in some claim R, then that agent's betting behaviour will typically be unstable across all bets of the form D, where $m < v < n$ (assuming as usual that the agent values only money, and values money linearly). And of course an agent will display similarly unstable betting behaviour across all bets which the agent would deem relevantly equivalent. With this in mind, we can now give a strengthened version of the two bets argument against imprecise probabilities. The challenge is to produce a rule that guides rational action, with the following features:

1. The rule respects the fact that where probabilities are imprecise, betting behaviour is typically unstable, in the sense described above. To fail to respect this would be to '[rob] unsharp probabilities of their point'—which is the accusation that Elga reasonably levels against the mid-point view (Elga 2010, p. 6).
2. The rule entails that an agent in the two-bets scenario[18] who rejects both bets is irrational.

We can see that the prevalent objection fails to meet this strengthened challenge. To ensure that the second requirement is met, the proponents of this view claim that the agent if rational will be acting in line with some sophisticated decision rule—such as maximin. And the idea is not merely that the decision rule is modelling the agent's behaviour: the agent's behaviour is literally described by the model. Further, the agent's doxastic state and preferences are stable and known, so the agent is able to accurately predict his or her own future behaviour. All of this is required if the rule is to entail that an agent who rejects both bets is irrational. But in meeting this second requirement, the proponents violate the first requirement: the solution mischaracterizes the phenomenon that it was designed to defend.

We can see that the permissive choice rules obviously violate the second requirement. And so we are left with the midpoint rule, the narrowing proposal, the sequence proposal and the planning proposal. Elga gives independent reasons for rejecting each of these rules, but it is worth noting here how each rule fares with respect to requirements 1 and 2 above. The midpoint rule clearly violates requirement 1. The narrowing proposal and the sequence proposal do not strictly violate requirement 1, but they do violate a natural extension of this requirement, which I turn to now.

I defined 'unstable betting behaviour' as betting behaviour that varies across close possible worlds, where the agent has the same evidence and preferences. Where betting behaviour varies in this way across close possible worlds, I claim that typically it will similarly vary within the same possible world across time—provided that the (relevant) evidence and preferences are unchanged. To illustrate the point here, suppose as usual that your probability in R is the range [0.1–0.8], that you value money linearly, and so on. At ten o'clock a bookie approaches you with two bets of the form D, one ($D_{0.5}$) with $v = 0.5$, and one ($D_{0.51}$) with $v = 0.51$: both bets are offered simultaneously, and you are free to accept or reject either or both of them. You decide which bets to accept, on a whim. Then, five minutes later, the bookie comes back to you full of apologies: the betting slips have blown away, and all bets made at 10 o'clock are called off. You are offered bets $D_{0.5}$ and $D_{0.51}$ once again, and you're free to give the same or a different answer to your earlier response—the bookie cannot even remember it. I claim that you might accept both bets at ten o'clock, but reject both at five past ten—or vice versa. In other words, your betting behaviour is unstable across time. I claim that where probabilities are imprecise, betting behaviour (even the betting behaviour of rational agents) exhibits a similar same sort of instability across time as it does across possible worlds.

We can then extend requirement 1 to include instability across time as well as across close possible worlds. If we do so, then we can see that the narrowing proposal violates this requirement. To see this, consider that according to the narrowing proposal, an agent who accepted bet $D_{0.51}$ at 10 o'clock (and so we can assume also accepted bet $D_{0.5}$) would thereafter (at least, until (s)he gained some new evidence) have in the set that represents his or her doxastic state only probability functions relative to which the expected utility of accepting bet $D_{0.51}$ is greater than—or at least equal to—the expected utility of rejecting bet $D_{0.51}$. Thus relative to every probability function in the set that represents the agent's doxastic set from 10 o'clock onwards, the probability assigned to R is greater than or equal to 0.51. It would be irrational, then, for the agent at five past ten to reject bet $D_{0.5}$—when relative to every probability function in the set representing the agent's doxastic state, accepting bet $D_{0.5}$ has a higher expected utility than rejecting it.

Do the sequence and planning proposals similarly violate the first requirement, once it is extended to cover temporal as well as modal instability? Here the proponent of the sequence proposal might make the following move: not just *any* series of actions counts as a 'sequence'. For a series of actions to count as a sequence, some special features must hold: e.g. perhaps the agent must suspect that (s)he will carry out this sequence, or perhaps (s)he must see the individual actions as in some sense part of a coherent larger action.[19] In general, where probabilities are imprecise, betting behaviour is unstable as described above across both possible worlds and time: thus requirement 1 is satisfied. However in the special case where a series of actions qualifies as a 'sequence', the agent (if rational) will not exhibit this sort of instability across time (though (s)he will still exhibit it across possible worlds). Instead, there will be at least one probability function in the set representing the agent's doxastic state, relative to which *every* action taken in the series maximizes expected utility. The series of bets in the two-bets scenario qualifies as a sequence: that is why a rational agent will not reject both bets. The scenario described above involving bets $D_{0.5}$ and $D_{0.51}$ is *not* a sequence, and this explains why a rational agent might accept both bets at the earlier time and then reject both at the later time. Thus the sequence proposal may be able to satisfy both requirements 1 and 2.

The proponent of the planning proposal can make a similar move. Rational agents are sometimes required to make plans and stick to them—but they are not required to plan for every possible eventuality. The two-bets scenario is a case where a rational agent *is* required to make a plan and stick to it. The plan must maximize expected utility relative to some probability function in the set that represents the agent's doxastic state—and no such plan will advocate rejecting both bets. This explains why it would be irrational to reject both bets in the two-bets scenario. However the example involving $D_{0.5}$ and $D_{0.51}$ is not a case where the agent is required to make a plan, and so the agent can be rational and yet exhibit unstable betting behaviour in this scenario. In general, agents with imprecise probabilities typically exhibit unstable betting behaviour across close possible worlds, and also across time—unless the case is one where the agent is rationally required to make a plan and stick to it. Thus the planning proposal may be able to satisfy both requirements 1 and 2.

Of course, the sequence proposal and the planning proposal still face the problems that Elga raises against them (Elga 2010). But I note here that they seem to work as well as responses to our strengthened version of the two-bets argument as they do to the original two-bets argument. We can see clearly now that the real challenge is to explain why it is that where subjective probabilities are (intuitively) imprecise, betting behaviour is typically unstable—but snaps into line in the two-bets scenario. Both the sequence proposal and the planning proposal have the right sort of shape to respond to this challenge. Whether they are satisfactory depends on whether Elga's objections to these theories can be met, and the proponents of these solutions have arguments in defense of their positions which I do not pursue here. My aim has been to show that the sequence and planning proposals are as applicable to the strengthened version of the two-bets argument as they are to the original two bets argument—unlike the prevalent response that I have attacked in this paper.

Section 6: Should Subjective Probabilities Be Sharp?

I end by considering Elga's own view—that rational agents have only sharp probabilities. Does this work well as a response to the two-bets scenario?

There are two ways that we might cash out the claim that 'subjective probabilities should be sharp'. We might claim that rational agents have sharp probabilities, without implying that this requires that betting behaviour should be stable. Or we might claim that rational agents have sharp probabilities, and further that sharp probabilities entail stable betting behaviour—and so conclude that rational agents always exhibit stable betting behaviour.[20] I consider each of these options in turn.

First of all, let us consider the idea that rational agents might have sharp probabilities even while exhibiting unstable betting behaviour. This might seem like an odd idea—for what could determine that an agent had some particular precise probability function while his or her betting behaviour is unstable? However there are a number of ways this idea could be cashed out. We could say that the precise probability is fixed somehow by the pattern of betting behaviour: perhaps the probability function is some sort of average of the agent's actual and dispositional betting behaviour, or perhaps the probability function supervenes on the betting behaviour in 'an unsurveyably chaotic way' (Williamson 1994, p. 209). We could alternatively say that the agent's subjective probability is a functional state—and perhaps it will somehow turn out that this functional state is indeed a relation to a precise probability function. Or we could say that the precise probability function of an agent is fixed in some way by his or her brain state.[21] Or perhaps we could say that the whole idea of an agent's having a probability function should not be taken too literally: it is merely a way of *modelling* the agent's doxastic state, and so even an agent with unstable betting behaviour gets modelled by a precise probability function. Any of these ideas—and there may be others—might be used to cash out the idea that an agent can have a sharp probability function even while his or her betting behaviour is unstable.

The claim that rational agents have sharp probability functions—when not saddled with the assumption that sharp probability functions entail stable betting behaviour—may be easy enough to swallow. On this view we can maintain that rational agents exhibit unstable betting behaviour, even while granting that rational agents have sharp probability functions. The problem however is that this view cannot handle the two-bets scenario. If rational agents have sharp probabilities, but sharp probabilities are no guarantee of stable betting behaviour, then (for all that has been said) rational agents may exhibit unstable betting behaviour. The fact that an agent who rejects both bets in the two-bets scenario is classed as irrational is then left unexplained. Why shouldn't a rational agent reject both bets in this scenario, given that rational agents can in general exhibit unstable betting behaviour? Thus merely stipulating that rational agents have sharp probabilities is not an adequate response to the two-bets scenario, if having sharp probabilities does not place any particular restriction on an agent's betting behaviour.

Let us then turn to the second idea—that rational agents have sharp probabilities, and that this entails that rational agents always have stable betting behaviour. This does indeed give us a response to the two-bets scenario, for an agent whose betting behaviour is always stable would not reject both bets. I expect that this is the position that Elga is arguing for. However it is worth noting how radical this position is. It is obvious enough that everyone—including intelligent, sane people such as yourself—exhibit unstable betting behaviour.[22] I have argued that unstable betting behaviour is a typical mark of those common cases where intuitively we think that an agent's probability function is not sharp. Elga's claim is that only irrational agents exhibit this sort of behaviour—and so it follows that we are all irrational in virtue of this behaviour. This is a counterintuitive claim—it is a mass 'irrationality theory', as unappealing as many maligned error theories. A natural way to try to soften the blow is to claim that 'irrationality' here is meant in some sort of technical sense. The idea is not that we are all just being stupid: rather, we are merely falling short of some notion of *ideal* rationally. The claim that we are not *perfectly* rational is something that Bayesians have long got used to. And of course Elga can coherently claim that an ideally rational agent will not exhibit unstable betting behaviour, and thus will not reject both bets offered in the two-bets scenario. The problem is that this does not explain why *you* would not reject both bets offered in the two-bets scenario. For recall that in the two-bets scenario we did not need to switch to some technical notion of ideal rationality to establish that it would be irrational to reject both bets. It's just an obviously stupid thing to do, and not something that *you* would do: it's irrational in the everyday sense. To argue that an ideally rational agent never exhibits unstable betting behaviour and therefore would not reject both bets does not account for the fact that you, and in general agents who are rational in the everyday sense would also not reject both bets. If Elga's claim that rational agents have sharp probabilities—and so have stable betting behaviour—is supposed to explain why it is irrational in the everyday sense to reject both bets, then the sense of 'rationality' had better be the same throughout. Thus the claim would need to be that all agents who are rational in the everyday sense have sharp probabilities, and so exhibit only stable betting behaviour. As we are surrounded by unstable betting behaviour and exhibit it ourselves, this claim would commit us to the radical view that irrationality is all around us where we had never suspected it—and not just technical irrationality either, but good old-fashioned stupidity.

Section 7: Conclusion

The two-bets argument is a challenge to the advocate of imprecise probabilities. I have strengthened this challenge, in the light of my claim that unstable betting behaviour is a typical mark of the cases that motivate imprecise probability theorists. A range of popular recent responses to the argument do not succeed against this strengthened version, and we are left with the sequence proposal and the planning proposal as serious contenders. A cost of accepting either of these proposals is the difficulty of handling the problems that Elga raises against them. Elga's own response to the scenario comes with its own cost: it entails mass irrationality of the everyday sort. There is no easy solution, then, to the two-bets argument.

Notes

[1] This paper has its roots in the Managing Severe Uncertainty Group at the LSE (supported by the British Academy), and I continued to work on this while supported by the Leverhulme Trust. Thanks to Katie Steele and Alex Voorhoeve for discussion on this topic, and thanks also for the many invaluable questions and comments from audiences at the Birmingham Philosophy Society, the Vienna Forum for Analytic Philosophy, the Franco-Swedish Program for Philosophy and Economics, and the Bristol-Groningen Conference in Formal Epistemology.

[2] You might produce the number '0.5', and feel that this is *not* an arbitrary choice because it lies exactly between 0 and 1. To see a problem with this response, consider the claim (R^+) that it will rain at midday *and* at 11.50am in London on New Year's day 2025. Your probability in R^+ will presumably be lower than your probability in R, so you can't have a probability of 0.5 in both claims. What, then, are your probabilities in these two claims?

[3] E.g. in Hammond (1988), Seidenfeld (1988), McClenen (1990), Levi (1991), Weatherson (2008).

[4] Many thanks to an anonymous reviewer for this point.

[5] The objection that I focus on is raised by Bradley and Steele (2014), but it is not their ultimate objection to Elga's argument.

[6] Moss (2015) has argued that we can construct formally similar cases where we would *not* intuitively judge that an agent who rejected both bets (or the equivalent in her scenarios) would be irrational. She argues that these cases are the ones where we judge that the agent 'genuinely changes her mind between bets' (Moss 2015, p. 675), and there are some features that typically obtain in these cases—such as a significant amount of time passing between the two bets being offered, and evidence of 'psychological effort' (Moss 2015, p. 676). I find this persuasive, and grant Moss's point. The two-bets scenario can still be constructed, however: we simply build into the scenario (as Elga did) that the bets are offered in close succession, and we ensure that the claim over which the bet is made is hardly worth agonizing over. In other words, we construct the scenario in such a way that none of the features typical of cases where an agent 'genuinely changes her mind' obtain. In this scenario—even by Moss's own lights—an agent who rejected both bets would be intuitively classed as irrational.

[7] There are cases (e.g. Newcomb's paradox) in which it is not clear that maximizing expected utility is the rational strategy, but we can set up the two-bets scenario to ensure that it does share the special features of these cases.

[8] Michael Bratman argues from considerations of 'self-governance' that intentions have 'a kind of default status'—meaning that a rational agent will retain an intention once formed for as long as there is no change in the agent's assessment of the reasons for forming that intention (Bratman 2012, p. 79). Other theorists ('resolute choice theorists') have also put forward arguments for the claim that rational agents are (all else being equal) required to follow their plans (e.g. McClenen 1990, Gauthier 1986). I do not delve into this debate here.

[9] This sort of principle is defended by Brian Weatherson, who writes: 'the Caprice rule [which is similar to a permissive choice rule] must be expressed in terms of the reasonableness of *sets* of decisions' (Weatherson 2008, p. 11, emphasis added). This is endorsed by Robert Williams (Williams 2014, p. 15), and James Joyce's response to Elga might also be classed as a version of the sequence proposal (Joyce 2011, pp. 317–318).

[10] In particular, there is no relevant difference in the agent's doxastic state (the agent will presumably remember being offered and refusing bet A in one scenario but not in the other, but the agent's probability in R will be the same across the two scenarios), and there is no relevant difference in the agent's preferences.

[11] Elga discusses maximin in a footnote (Elga 2010, p. 5), where he claims that maximin is effectively a permissive choice rule. However in his errata (Elga 2012, p. 1) Elga corrects this claim, and considers a version of the prevalent objection that we are concerned with here.

[12] This holds provided that certain other conditions—carefully detailed by Chandler (2014, p. 12)—are met.

[13] This needs to be qualified: in cases where two or more possible actions have the same maximum minimum utility, then maximin does not determine which action a rational agent will perform—for all actions that have maximum minimum utility are rationally permissible. This sort of exception holds for each of the relevant rules (e.g. for maximax, Hurwicz's rule, and so on). The general strategy for responding to Elga that I focus on is unaffected by this limited sort of indeterminacy.

[14] I hold off from saying that this sort of instability is a *necessary* ingredient for cases where (intuitively) agents have imprecise subjective probabilities, because there may be cases where an agent deliberately regulates her choices in some way whenever evidence (plus preferences) does not dictate her choice. For example, we can imagine an agent who accepts or rejects bets as her evidence and preferences rationally dictate, and in un-dictated cases always accepts. Of course, her decision to accept must be arbitrary—i.e. not dictated by her evidence and preferences: we shouldn't imagine here an agent who on the whole enjoys accepting offers more than rejecting them (perhaps out of politeness)—because this preference should have played its role along with the agent's evidence in rationally dictating her choice. We must instead imagine an agent who has chosen arbitrarily to follow this rule. If such a case is possible, then we have a case where the agent's betting behaviour is arbitrary—and not dictated by her evidence and preferences—but nevertheless is stable. It is because of the possibility of this sort of case that I say that unstable betting behaviour is a *typical* (rather than *necessary*) feature of cases where intuitively we claim that an agent has an imprecise subjective probability.

[15] We are still assuming here that you value only money, and value money linearly.

[16] This is a case of 'inexact knowledge' (Williamson 1994, pp. 216–234). If we suppose that there are in fact w sweets in the jar, then you know that there are at least v sweets only if v is smaller than w by some margin (the 'margin for error'). Williamson argues that you do not know that there are at least v sweets where v is within the margin for error, because there are close possible worlds where your belief is false. There are (at least) two ways that your belief might be false in a close possible world: you might have a belief with the same content as in the actual world, even though the number of sweets in the jar is different; or you might have a belief with a different content (you might easily have believed otherwise), despite the number of sweets in the jar being the same. Williamson discusses both sorts of cases, but it is the second sort of case that is most relevant to this paper. To bring out the relevance, replace 'believes' with 'guesses': where v is within the margin for error, an agent who guesses that there are at least v sweets in the jar could easily have guessed otherwise—and so (in other words) there will be close possible worlds where the agent does guess otherwise.

[17] Perhaps you could discover all the relevant factors: I don't mean to argue that it is in principle impossible to predict your own or others' betting behaviour. The key point is just that in every day cases of the relevant sort you typically cannot make these predictions, perhaps partly because you just do not have a brain scanner to hand. What holds of the everyday cases holds of the two-bets scenario too: you do not have a brain scanner to hand in that scenario either. If you *are* able to make the relevant predictions in the two-bets scenario but not in everyday cases, this demands an explanation.

[18] Here as before we construct the two-bets scenario in such a way that even on Moss's view an agent who rejected both bets would be deemed irrational. See footnote 6.

[19] Jonathan Weisberg suggests that Elga slips between describing the agent's response to the bookie's offers of bets A and B as a single action, and as two separate actions (Weisberg 2015, pp. 824–825). Analogously, a proponent of the sequence proposal might claim that Elga's argument plays on our uncertainty as to whether the agent's response to the pair of bets qualifies as a sequence.

[20] I note that even where an agent (intuitively) has sharp probabilities, that agent will still display unstable betting behaviour over bets that (s)he considers fair. E.g. an agent with a precise probability of y in R will display unstable betting behaviour over a bet of the form D with $v=y$. The real contrast then is between betting behaviour that is unstable over bets of the form D for a range of values v, and betting behaviour that is unstable over bets of the form D only for a unique value v: for simplicity however in this section I refer to these as 'unstable' and 'stable' betting behaviours respectively.

[21] As Ramsey writes: 'it is, I suppose, conceivable that degrees of belief could be measured by a psychogalvanometer or some such instrument . . . ' (Ramsey 1926, p. 161).

[22] Indeed, we exhibit unstable betting behaviour across bets of the form D for a *range* of values *v*: see footnote 20.

Bibliography

Arrow, K. J. & Hurwicz, L. (1972). An Optimality Criterion for Decision Making under Ignorance. In C. F. Carter & J. L. Ford, *Uncertainty and Expectations in Economics: Essays in Honour of G.L.S. Shackle*. Oxford: Basil Blackwell.

Bradley, R. (2009). Revising Incomplete Attitudes. *Synthese*, *171* (2), 235–256.

Bradley, S. & Steele, K. (2014). Should Subjective Probabilities be Sharp? *Episteme*, *11* (3), 277–289.

Bratman, M. (2012). Time, Rationality and Self-Governance. *Philosophical Issues*, *22* (1), 73–88.

Chandler, J. (2014). Subjective Probabilities Need Not be Sharp. *Erkenntnis*, *79* (6), 1273–1286.

Elga, A. (2012). Errata for Subjective Probabilities should be Sharp. http://www.princeton.edu/~adame/papers/sharp/sharp-errata.pdf.

Elga, A. (2010). Subjective Probabilities Should be Sharp. *Philosophers' Imprint*, *10* (5), 1–11.

Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics*, *75* (4), 643–669.

Gärdenfors, P. & Sahlin, N. E. (1982). Unreliable Probabilities, Risk Taking and Decision Making. *Synthese*, *53*, 361–386.

Gauthier, D. (1986). *Morals by Agreement*. Oxford: Clarendon Press.

Hammond, P. (1988). Orderly Decision Theory. *Economics and Philosophy*, *4*, 292–297.

Jeffrey, R. (1983). Bayesianism with a Human Face. In J. Earman, *Testing Scientific Theories*. Minneapolis: University of Minnesota Press.

Joyce, J. (2011). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, *24* (1), 281–323.

Levi, I. (1991). Consequentialism and Sequential Choice. In M. Bacharach & S. Hurley, *Foundations of Decision Theory*. Oxford: Basil Blackwell.

Levi, I. (1974). On Indeterminate Probabilities. *Journal of Philosophy*, *71*, 391–418.

McClenen, F. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge: Cambridge University Press.

Moss, S. (2015). Credal Dilemmas. *Noûs*, *49* (4), 665–683.

Ramsey, F. (1926). Truth and Probability. In F. Ramsey, *The Foundations of Mathematics and Other Logical Essays*. New York: Harcourt, Brace and Company.

Sahlin, N. E. & Weirich, P. (2012). Unsharp Sharpness. *Manuscript*.

Satia, J. & Lave, R. (1973). Markovian Decision Processes with Uncertain Transition. *Operations Research*, *21* (3), 728–740.

Seidenfeld, T. (1988). Decision Theory without 'Independence' or without 'Ordering'. *Economics and Philosophy*, *4*, 267–290.

van Fraassen, B. (1990). Figures in a Probability Landscape. In J. M. Dunn & A. Gupta, *Truth or Consequences: Essays in Honor of Nuel Belnap*.

Weatherson, B. (2008). Decision Making with Imprecise Probabilities. *Manuscript*.

Weisberg, J. (2015). You've Come a Long Way, Bayesians. *Journal of Philosophical Logic*, *44* (6), 817–834.

Williams, R. (2014). Decision-Making Under Indeterminacy. *Philosophers' Imprint*, *14* (4).

Williamson, T. (1994). *Vagueness*. London: Routledge.