Philosophers' Imprint

volume 10, no. 5
may 2010

Subjective Probabilities should be Sharp

Adam Elga
Princeton University

Abstract

Many have claimed that unspecific evidence sometimes demands unsharp, indeterminate, imprecise, vague, or interval-valued probabilities. Against this, a variant of the diachronic Dutch Book argument shows that perfectly rational agents always have perfectly sharp probabilities.

1. Introduction

Sometimes one's evidence for a proposition is *sharp*. For example: You've tossed a biased coin thousands of times. 83% of the tosses landed heads, and no pattern has appeared even though you've done a battery of statistical tests. Then it is clear that your confidence that the next toss will land heads should be very close to 83%.

Sometimes one's evidence for a proposition is *sparse but with a clear upshot*. For example: You have very little evidence as to whether the number of humans born in 1984 was even. But it is clear that you should be very near to 50% confident in this claim.

But sometimes one's evidence for a proposition is *sparse and unspecific*. For example: A stranger approaches you on the street and starts pulling out objects from a bag. The first three objects he pulls out are a regular-sized tube of toothpaste, a live jellyfish, and a travel-sized tube of toothpaste. To what degree should you believe that the next object he pulls out will be another tube of toothpaste?

The answer is not clear. The contents of the bag are clearly bizarre. You have no theory of "what insane people on the street are likely to carry in their bags," nor have you encountered any particularly relevant statistics about this. The situation doesn't have any obvious symmetries, so principles of indifference seem to be of no help.

Should your probability be 54%?

91%?

18%?

It is very natural in such cases to say: You shouldn't have *any* very precise degree of confidence in the claim that the next object will be toothpaste. It is very natural to say: Your degree of belief should be *indeterminate* or *vague* or *interval-valued*. On this way of thinking, an appropriate response to this evidence would be a degree of confidence represented not by a single number, but rather by a range of numbers. The idea is that your probability that the next object is toothpaste should not equal 54%, 91%, 18%, or any other particular number. Instead it should span an interval of values, such as [10%, 80%].[1]

The toothpaste-in-the-bag example is artificial, but many realistic examples have been proposed. What is your confidence that "there will be a nuclear attack on an American city this century"?[2] What is your state of opinion concerning "the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth owners in the social system in [40 years]"?[3] It is tempting to agree with J. M. Keynes that "About these matters there is no scientific basis on which to form any calculable probability whatever"[4] and to think that the problem isn't just that our computers aren't fast enough.

Let me emphasize: The idea is not that some computational or representational limitation prevents you from having a definite probability. Give an agent access to exactly your evidence relevant to the toothpaste claim, or, say, the claim that there is a God. Give her all the computers, representational tools, brain upgrades, etc. that you like. *Still* it seems as though the agent would go wrong to have any very precise degree of belief in the relevant claim. The evidence itself just seems too unspecific to warrant a precise degree of belief, no matter how smart the person evaluating the evidence is.

James M. Joyce expresses the point particularly clearly:

> As sophisticated Bayesians like Isaac Levi (1980), Richard Jeffrey (1983), Mark Kaplan (1996), have long recognized, the proper response to symmetrically ambiguous or incomplete evidence is not to assign probabilities symmetrically, but to refrain from assigning precise probabilities at all. Indefiniteness in the evidence is reflected not in the values of any single credence function, but in the spread of values across the family of all credence functions that the evidence does not exclude.... It is not just that sharp degrees of belief are psychologically unrealistic (though they are). Imprecise credences have a clear epistemological motivation: they are the proper response to unspecific evidence.[5]

Other authors agree.[6] Here is Isaac Levi:

> it is sometimes rational to make no determinate probability judgment and, indeed, to make maximally indeterminate judgments.... [Doing so] may derive from a very clear and cool judgment that on the basis of the available evidence, making a numerically determinate judgment would be unwarranted and arbitrary.[7]

Here is Scott Sturgeon:

> When evidence is essentially sharp, it warrants a sharp or exact attitude; when evidence is essentially fuzzy—as it is most of the time—it warrants at best a fuzzy attitude.[8]

1. Perhaps the interval itself should have vague boundaries, as is suggested in Sturgeon (2008). Or perhaps the spread should be represented not by having an interval range, but rather by having your state of belief represented by a set of probability functions rather than a single one, as in Levi (1980), Jeffrey (1983), Joyce (2005), Kaplan (1996) and van Fraassen (2006). In fact the set-based framework is superior, but the details won't matter for the argument of this paper. So in the main text I talk in terms of interval-valued probability functions. For a helpful review of representation theorems that involve imprecise probabilities or utilities, see Seidenfeld et al. (1990).
2. Weatherson (2005).
3. Keynes (1973), as cited in Weatherson (2005).
4. Keynes (1973).
5. Joyce (2005, 171).
6. In addition to the works listed below, Kaplan (1996) also rejects the idea that rationality requires sharp probabilities, naming it "the sin of false precision." See Kaplan (1996, 23–31) and also Kaplan (2009).
7. Levi (1985, 396).
8. Sturgeon (2008, 27).

Here is Peter Walley:

> If there is little evidence concerning [a claim,] then beliefs about [that claim] should be indeterminate, and probability models imprecise, to reflect the lack of information. We regard this as the most important source of imprecision.[9]

These authors all agree that one's evidence can make it downright unreasonable to have sharp degrees of belief. The evidence itself may call for unsharp degrees of belief, and this has nothing to do with computational or representational limitations of the believer. Let me write down a very cautious version of this claim:

> UNSHARP: It is consistent with perfect rationality that one have unsharp degrees of belief.

This claim is initially quite plausible. But it is false. What's true is:

> SHARP: Perfect rationality requires one to have sharp degrees of belief.

2. Sharp does not entail Uniqueness

Before a defense of SHARP, one clarification: SHARP does not say that for every batch of evidence, only *one* subjective probability function is rationally permissible. In the lingo: It does not entail "Uniqueness".[10] SHARP says that perfect rationality requires *that one's probability function be sharp*. It is compatible with SHARP that for certain batches of evidence, there is more than one probability function it is rationally permissible to have on the basis of that evidence. SHARP just demands that if that is so, each such permissible function is perfectly sharp.[11]

9. Walley (1991, 212-3).
10. Feldman (2007), White (2005).
11. There may well be difficulties with accepting SHARP while denying UNIQUENESS. But I will not press any such difficulties here. Thanks to Susanna Rinard and John Collins for pressing me on this point.

3. How do unsharp probabilities constrain rational choice?

So: do any perfectly rational agents have unsharp probabilities? I will argue that the answer is "no." My reason is that there is no plausible account of how unsharp probabilities constrain the reasonable choices of such agents. Let me start by explaining what it means to ask for such an account.

According to the standard story (expected utility theory), a perfectly rational agent's degrees of belief are represented by a probability function. Such an agent always performs an action that has maximal expected utility. In this paper, we may restrict ourselves to the simplest case, in which the agent's utility scale is "linear with dollars."[12] That just means that her expected utility gain for a simple betting arrangement[13] equals a probability-weighted average of the various monetary gains that the arrangement could yield. For example, if an arrangement yields a gain of $50 if it rains, and a loss of $4 otherwise, and the agent's probability for rain is 20%, then her expected gain for that arrangement is $(20\%)(50) + (80\%)(-4) = 6.8$. In this example, the expected gain for accepting the arrangement (6.8) is greater than the expected gain for rejecting it (0). So if those are her only options, the agent will accept the arrangement.

That's the standard story, which presupposes that perfectly rational agents always have perfectly precise probability functions. Anyone who claims that such agents can have unsharp probability functions owes a corresponding account of how unsharp probabilities constrain rational action.

12. The assumption of linear utility is made purely for convenience. All of of the examples could be modified to work for any nontrivial utility scale.
13. A betting arrangement is simple when the agent counts the proposition being bet on as independent of whether she accepts the bet. For example, bets by ordinary people on tomorrow's weather are simple, as are bets on who will win the world series. A bet on the world series may not be simple if it is made by someone who will play in the series. For in that case, the player may count his betting choices as relevant to who will win. I assume that all of the betting arrangements mentioned in the paper are simple. I also ignore subtleties associated with causal decision theory, since they are not relevant to present concerns.

Now for an argument that no such account is acceptable.

4. A great series of bets

Let H be some particular proposition. For example, H might be the proposition that it will rain tomorrow. I'm going to offer you a great series of bets on H:

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.

Since the bets will be offered in quick succession, you're sure that as regards the truth of H, your state of mind will remain exactly the same throughout the whole bet-offering process. That means: You won't get any new evidence relevant to H. You won't lose any evidence regarding H. You won't have any change of mind about how to interpret, assess, or respond to your existing evidence. You won't have any religious revelations, conversion experiences, or other kind of change in your opinion regarding whether H is true. Just the regular old passage of time, and me offering you this great sequence of bets.

The bets are great because if you accept both of them, you'll be sure to win $15 on one, lose $10 on the other, and so gain $5 overall. In other words, bets A and B together guarantee a sure gain, just as a Dutch Book guarantees a sure loss.[14]

Now, it is not rationally required that you accept both of the bets.[15] For example, you might be so confident in H that accepting Bet B alone is even more attractive than accepting both bets. Or you might be so doubtful of H that accepting Bet A alone is even more attractive than accepting both bets.

14. On such arrangements, which are in the literature sometimes termed "arbitrage opportunities," see Schervish et al. (1998, 142) and Hájek (2005).
15. Cf. Kyburg (1978, fn. 2). Thanks here to Alan Hájek and Teddy Seidenfeld.

But it *is* rationally required that you accept at least one of the bets. For rejecting both bets is worse for you, no matter what, than accepting both bets. And you can see that in advance. So no matter what you think about H, it doesn't make sense to reject both bets.

That claim is the main premise of this paper, and so deserves emphasizing:

> Any perfectly rational agent who is sequentially offered bets A and B in the above circumstances (full disclosure in advance about the whole setup, no change of belief in H during the whole process, utilities linear in dollars) will accept at least one of the bets.

Those who think rational agents always have sharp probability functions can easily accommodate this claim. I will now argue that defenders of UNSHARP cannot accommodate it at all.

5. Permissive choice rules

Let the setup be the same as before: you are informed in advance that you will be offered Bet A followed by Bet B and that your state of opinion regarding H will remain unchanged throughout the whole process. Assume that UNSHARP is right, and that in this case your probability for rain tomorrow should be a wide interval. For definiteness, suppose that you are perfectly rational and that you have $P(H) = [10\%, 80\%]$. (The wide interval in this example is chosen purely for convenience. Given any agent who assigns imprecise probability to at least one proposition H, there exist a pair of bets capable of playing the role that bets A and B play in the argument below.)

Given your imprecise degree of belief in H, how are you rationally required to respond to the above bets?

First consider Bet A. It is clear how a rational agent with sharp probabilities would evaluate Bet A: She would accept it if her $P(H)$

were less than 60%. She would reject it if her $P(H)$ were greater than 60%. And she would be rationally permitted to either accept or reject it if her $P(H)$ were exactly 60%.

But what about you? Your $P(H)$ is an interval that extends from well below 60% to well above 60%. One natural thing to say is that rationality counts the bet as optional for you—it neither requires you to accept it nor to requires you to reject it. And that thought is very much in the spirit of UNSHARP. If your evidence is so unspecific as to demand a widely spread-out probability function, it is natural that the requirements of rationality be correspondingly spread out. It is natural that there be a whole range of bets such that for each one, it is rationally permitted that you accept it and also rationally permitted that you reject it.

It is natural to say the same thing about Bet B: that when it is offered, it too is rationally optional.

Finally, it is natural to conclude that a perfectly rational agent may reject both bets.

But that is absurd, so this very natural proposal is of no use to the defender of UNSHARP. More generally, this reasoning shows that if UNSHARP is true, none of a whole range of choice rules are acceptable. That range includes rules given by Isaac Levi,[16] Peter Walley,[17] I. J. Good,[18] Teddy Seidenfeld, [19] Marco Wolfenson and Terrence L. Fine,[20] Peter Gärdenfors and Nils-Eric Sahlin,[21] and Itzhak Gilboa and David Schmeidler.[22] These rules all entail that in the above case it can be rationally permissible to reject both bets.

16. Levi (1980, 1985). Levi fairly and squarely faces the sort of phenomenon exhibited by bets A and B. He recognizes the need to say how indeterminate probabilities constrain rational decision, and gives a rule that does so. The rule assumes that an agent's belief state is represented by a set of probability functions. The rule says that an act is rational iff it (1) maximizes expected utility according to at least one of the functions in the agent's set, and (2) meets some additional conditions which depend on the agent's value judgments (Levi 1980, 162). Given certain value judgments that the rule counts as rationally permissible (for example, ones which make clause (2) impose no extra conditions), the theory counts it as rationally permissible to reject both bets A and B in the above situation. Levi both recognizes and embraces this consequence of his theory. Addressing a similar case, he writes: "According to the permissibility interpretation, several distributions are permissible in both decision problems. Some favor one option and some favor another. There is not the slighted obligation on the rational agent to choose in a way that optimizes relative to the same probability distribution in the two distinct hypothetical contexts." (Levi 1985, 395) Levi has two main motivations for embracing this consequence. First, he sees it as an unavoidable consequence of thinking that perfectly rational agents may have indeterminate probability judgments, which he takes to be overwhelmingly plausible. Second, he takes it that denying the consequence imposes "an excessive rigidity [over time] in our judgments of probability" (Levi 1987, 208).

17. Walley (1991, Section 5.6.6) describes several rules for how unsharp probabilities determine rationally permissible choices. All of these rules entail that it is rationally permissible to reject both bets A and B in the above situation.

18. Good (1952, 114).

19. See, for example, Seidenfeld (1984, fn 1): "I have heard it said that General George C. Marshall defined a coward as one who would not accept either the bet on E at odds of 1:2 or the bet on –E at odds of 1:2. Whether a 'spread in the odds' reflects cowardice I cannot say, but I am sure that cowards may be rational." Seidenfeld et al. (1990, 278) and Seidenfeld (1994, Section 2.1) anticipate the sort of phenomenon exhibited in the Bet A/Bet B situation, and distinguish sharply between sequential and non-sequential versions of a given decision problem. In particular, they impose no constraints that rule out that a rational agent might reject both bets in the above situation. Note that the recommendations of this theory depend on the "security rule" that the agent employs, and some choices of security rule will require the agent to accept at least one of the bets. Special thanks to Teddy Seidenfeld for detailed and helpful correspondence on this point.

20. Wolfenson and Fine (1982).

21. Gärdenfors and Sahlin (1982). In the framework of Gärdenfors and Sahlin (1982), unsharp belief states are represented by sets of probability distributions, each of which has been tagged by a number indicating a degree of "epistemic reliability." The decision rule is that one choose the option with "the largest minimal expected utility." (371) The minimal expected utility of an option is the minimum of the expected utilities assigned to that option by any of the sufficiently reliable probability functions that are members of one's probability set. In many circumstances, this rule entails that it can be rational to sequentially reject both Bet A and Bet B.

22. Gilboa and Schmeidler (1989). To be fair, nothing in Gärdenfors and Sahlin (1982) or Gilboa and Schmeidler (1989) requires that *perfectly* rational agents may have unsharp probabilities. Those theories may be understood as saying how a non-ideal agent's unsharp probabilities should constrain his actions. So understood, the arguments given here make no trouble for the theories.

One might insist that rejecting both bets in the above situation *is* consistent with perfect rationality. I don't find this plausible, but have no additional arguments. I can only ask you to vividly imagine a case in which an agent rejects both bets A and B. Keep in mind that this agent cares only about money (her utility scale is linear), that she is certain in advance what bets will be offered, and that she is informed in advance that her state of opinion on the bet proposition will remain absolutely unchanged throughout the process. I invite you to agree that this agent has exhibited a departure from perfect rationality.

6. Strict rules

The defender of UNSHARP owes a story explaining how unsharp probabilities constrain rational action. We've seen that some such stories wrongly entail that it is permissible to reject both bets in the above situation. What stories deliver the correct verdict, that it is irrational to reject both bets in the above situation? I can think of only a few kinds.

The first kind of story puts forward a very strict rule connecting unsharp beliefs to permissible betting odds. For example, consider the "midpoint rule," according to which agents should evaluate bets according to the midpoints of their probability intervals. According to that rule, an agent whose probability in H is the interval [10%, 80%] should evaluate bets on H in the same way as an agent whose probability in H is exactly 45%.

The midpoint rule (and other strict relatives of it), do indeed yield the verdict that it is impermissible to reject both bets. They yield that verdict because no rational agent with precise probabilities rejects both bets, and such rules require agents to bet just as if they had precise probabilities.

But such strict rules have another difficulty: they undermine the original motivation for UNSHARP. Let me explain.

Think back to the examples that initially motivated UNSHARP. Think of the toothpaste example. It initially seems implausible that rationality compels you to have a precise-to-the-millionth-decimal-place degree of belief that the next object out of the stranger's bag will be toothpaste. Likewise it initially seems implausible that rationality requires you to have a completely precise degree of belief that there is a God. There seems to be something about your evidence in these and other cases that fails to nail down an exact probability, and indeed rules out any particular exact probability.

But if that's so—if your evidence doesn't nail down an exact probability—it would be very strange if it *did* nail down a completely precise pattern of rational betting odds. And according to the midpoint rule, your evidence *does* nail down a completely precise pattern of rational betting odds. For your evidence nails down an interval, and the midpoint rule says that you should evaluate bets in the same way as someone whose probability is at the *exact* center of that interval.

For example, let H be the claim that the next object out of the stranger's bag is toothpaste. And suppose that you have $P(H) = [10\%,\ 80\%]$. Now consider the following ticket:

Worth $100 if *H* is true. Worth nothing otherwise.

According to the midpoint rule you should count that ticket as being worth exactly $45.000... . In other words, you should be willing to pay any amount up to $45 for it, and not a penny more. But if there is something fishy about rationality requiring you to have *exactly* probability 45% in H, then there is something just as fishy about rationality requiring you to value the H-ticket at *exactly* $45.

So the midpoint rule robs unsharp probabilities of their point. With the midpoint rule in place, interval-valued probabilities yield exact "point-valued" constraints on rational betting odds. But if it is always OK to have point-valued constraints on betting odds, there is no good reason for objecting to point-valued probabilities in the first place.

Bottom line: the midpoint rule undermines a main motivation for UNSHARP, and so cannot be used to buttress UNSHARP. The same goes for other strict rules.

7. Global rules

The state of play: we're considering rules that say how unsharp probabilities constrain rational action. We've considered permissive rules, according to which it can be rationally permissible to reject both Bet A and Bet B. These rules are too permissive, since it is irrational to reject both bets. We've considered strict rules, according to which unsharp probabilities completely determine sharp rational betting odds. Such rules are too strict, since they fail to do justice to the thought that motivates UNSHARP in the first place.

What a defender of UNSHARP needs is a rule that avoids both of the above difficulties. The rule should be strict enough to entail that it is irrational to reject both bets A and B. But it should be permissive enough to allow that sometimes, unsharp probabilities leave open a whole range of rationally permissible actions. In particular, it should allow that in a case in which Bet B is offered on its own, accepting Bet B is rationally optional. And it should motivate these conclusions in a natural way.

I can think of only three sorts of rules that fit the bill: rules according to which actions tend to *narrow* probability intervals, rules that appeal to the special role of *plans* in decision-making, and rules that appeal to *sequences of actions*.

I will now argue that none of these sorts of rules are acceptable.

8. The narrowing proposal

Start with the following proposal: When a rational agent with unsharp probabilities performs an action, her probabilities typically become more sharp or well-defined than they were before. They do so in such a way that her future actions will cohere with the action she just performed. For example, suppose that you have $P(H) = [10\%, 80\%]$. And suppose that you reject Bet A:

Bet A If H is true, you lose \$10. Otherwise you win \$15.

Then according to the present proposal—call it "NARROW"—when you reject Bet A, your $P(H)$ narrows to [60%, 80%]. Given this narrower interval, you will be inclined to accept Bet B if it should ever be offered to you. More generally, your intervals narrow in a way that prevents you from performing predictably inferior sequences of actions.[23]

What is attractive about NARROW is that it delivers the desired combination of strictness and permissiveness. It is strict enough to disallow sequences of actions that lead to predictably inferior outcomes. But it is permissive enough to allow that in many cases, unsharp probabilities license a range of rationally permissible actions. (For example, NARROW entails that in the Bet A/Bet B situation, rejecting both bets is irrational. That's because any rational agent who rejects Bet A will thereby sharpen her probabilities in such a way that she will accept Bet B. But the proposal also allows that in some cases in which Bet B is offered on its own, accepting Bet B is rationally optional. That's because when a rational agent is offered Bet B on its own, her probabilities might still be unsharp enough to make Bet B rationally optional.)

What is unattractive about the proposal is the way it requires perfectly rational agents to change their opinions without changes in their relevant evidence.

For example, let H be the proposition that it will rain today. Suppose that you are rational and that your probability for H is unsharp enough that each of the following actions are rationally permissible for you: (1) Wear an uncomfortable rain-poncho; (2) Wear a non-water-resistant suede jacket. According to NARROW, if you wear the rain-poncho you will become confident that it will rain, and if you wear the suede jacket you will become confident that it won't rain.

23. How exactly do your intervals narrow? The details are most simply filled in when an unsharp state of mind is represented not by an interval-valued probability function, but rather by a set of ordinary probability functions. Given that representation, NARROW might say: when one acts, one eliminates from one's set of probability functions those functions that do not endorse that action as rationally permissible. Since the details of the proposal don't matter for the objection I wish to raise, I suppress them in the main text and speak in terms of narrowing intervals.

Notice that in neither case will you gain or lose any evidence relevant to whether it will rain. For in each case, the only change in your evidence is this: you learn what jacket you choose to wear. And that news is not relevant to whether it will rain. In particular, you realize that your choice of jacket has no influence over the weather. And you realize that you don't have special rain-sensing powers that express themselves through inclinations to choose a jacket.

So according to NARROW, you are rationally required to change your opinion on whether it will rain even though your relevant evidence remains unchanged. But that's wrong. No perfectly rational agent is required to change her opinion on a subject matter when her relevant evidence remains unchanged. So NARROW is incorrect.

9. The planning proposal

NARROW says that when you choose to wear a rain-poncho, that should change your opinion regarding the rain. That is one way to guarantee that your subsequent rational choices will cohere with your choice to wear the poncho. But it is not the only way. One might instead say the following: when you choose to wear the poncho, you should make appropriate *plans* constraining your future choices, without at all changing your beliefs regarding the rain. In particular, you should plan to have your future choices cohere with your choice to wear the poncho. Then you should live according to your plans.

That thought motivates the following proposal (which was described but not endorsed in Dougherty (2007)):

> PLAN: Whenever a rational agent with unsharp credences performs an action, she simultaneously forms a plan governing her later actions. That plan requires her later actions to cohere with the action she just performed. If nothing unforeseen happens, she then follows through on her plan. In particular, in the Bet A/Bet B situation, whenever a rational agent rejects Bet A, she also simultaneously plans to accept Bet B. Later, when she is offered Bet B, she implements her plan.

Like NARROW, PLAN delivers the desired mix of strictness and permissiveness. It is strict enough to disallow sequences of actions that lead to predictably inferior outcomes. But it is permissive enough to allow that in many cases, unsharp probabilities license a range of rationally permissible actions.

(For example, PLAN entails that in the Bet A/Bet B situation, rejecting both bets is irrational. That's because any rational agent who rejects Bet A will thereby form a plan to accept Bet B, and will later follow through on that plan. But the proposal also allows that in a case in which Bet B is offered on its own, accepting Bet B is rationally optional. That's because when a rational agent is offered Bet B on its own, she need not have formed any previous plans that constrain her choice.)

To see the trouble with PLAN, consider Sally. Sally cares only about money (her utility scale is linear), and she has a highly unsharp degree of belief that it will rain. Contrast two situations. In the first situation, Sally rejects Bet A and then is offered Bet B. In the second situation, she is only offered Bet B.

According to PLAN, Sally is rationally permitted to reject Bet B in the second situation, but not in the first. But notice that in each situation, the monetary consequences of rejecting Bet B are exactly the same. And in each situation, the monetary consequences of accepting Bet B are exactly the same. And by assumption, Sally cares only about money. Furthermore, in each situation Sally's beliefs regarding the rain are exactly the same, and she knows just what the consequences of are of all of her choices.

So in each situation, the consequences of accepting Bet B and rejecting Bet B are exactly the same in every respect that Sally cares about.

To repeat: in the two situations, when Sally is deciding whether to accept Bet B she faces choices that are exactly the same in every respect that she cares about. So it can't be that rationality imposes different requirements on her in the two situations.[24] But according to

24. In making this inference, I have assumed that in each situation, Sally is

PLAN, rationality does impose different requirements on her in the two situations: it permits her to reject Bet B in the first situation, but not the second. So PLAN is incorrect.

It might be objected: "The two situations do differ in a way that matters: In the first situation but not the second, rejecting Bet B would violate Sally's *plan*. And it is irrational to violate one's plan without good reason."

In reply I ask: does Sally find it at all undesirable to break her plans? If she does—for example, if breaking her plans makes Sally feel guilty, or if it involves costly or unpleasant reconsideration—then the objection depends on misunderstanding the case. It was assumed that Sally cares only about money, and that reconsidering her plan causes Sally no discomfort.

If Sally doesn't find it at all undesirable to break her plans, then it is mysterious why it should count against Sally's doing something that it will break her plans.

One might posit a brute independent constraint on rationality: *Don't break plans!* But without further explanation, such a constraint is about as plausible as the constraint: *Don't break mirrors!* And that constraint isn't plausible at all. For either you find breaking mirrors undesirable in a given situation, or not. If you do, then there's some reason not to break them. If you don't, there isn't. And the same goes for plans.[25]

Moral: PLAN is incorrect.

certain just what consequences follow from her available choices.

25. Some theorists—*resolute choice theorists*—disagree. According to these theorists, in certain circumstances a rational agent is required to follow her plan, even when following that plan doesn't particularly serve her desires at the time. Such theorists appeal to circumstances in which plan-keepers—also known as *resolute choosers*—do better than others at satisfying their desires (Gauthier 1986, 1997, 1998, McClennen 1990, 2004). I can't hope to do justice here to the large and nuanced literature concerning resolute choice. For a convincing objection to resolute choice adapted from Smart's "rule-worship" objection to rule utilitarianism (Smart 1956, 348), see Bratman (1992, 8–10). For a reply on behalf of the resolute choice theorist, see DeHelian and McClennen (1993).

10. The sequence proposal

Turn now to one final proposal for how unsharp probabilities constrain rational action:[26]

> SEQUENCE: 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 also rationally permissible. But her performing the *sequence* of actions "reject-Bet-A-then-reject-Bet-B" was rationally impermissible.

SEQUENCE offers the same desirable mix of permissiveness and strictness as PLAN: It is strict enough to disallow sequences of actions that lead to predictably inferior outcomes. And it is permissive enough to allow that in many cases, unsharp probabilities license a range of rationally permissible actions.

To see why SEQUENCE is incorrect, consider again the 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 faces choices that are exactly the same in every respect she cares about. That is because the monetary consequences of accepting or refusing Bet B are the same in each situation, and Sally only cares about money.

So it must be that rationality imposes the same constraints on her in the two situations.

It might seem that SEQUENCE is consistent with this verdict. For SEQUENCE allows that in each situation the action "reject Bet B" is rationally permissible for Sally.

26. I adapt SEQUENCE from the "compound action" proposal offered in Hare (2007) and the principle "Caprice" defended in Weatherson (2008). A proposal in this spirit is also suggested in outline by Weirich (2001, 439–440).

But in fact, SEQUENCE entails that rationality imposes different requirements on Sally in the two situations. For according to SEQUENCE, the following is true: In order to be perfectly rational it is not enough to avoid irrational actions. One must also avoid irrational *sequences* of actions. In particular, Sally would be irrational if she were to reject Bet B in the second situation. For her doing so would complete the irrational sequence of actions "reject-Bet-A-then-reject-Bet-B." In contrast, her rejecting Bet B in the first situation would not complete any irrational sequences. So according to SEQUENCE, Sally's rejecting Bet B is consistent with her perfect rationality in the first situation, but not in the second.

Bottom line: SEQUENCE entails that rationality imposes different requirements on Sally in the two situations. But Sally can see that her choices in the two situations are alike in every respect that she cares about. So it must be that rationality imposes on her the same constraints in the two situations. So SEQUENCE is incorrect.

11. Conclusion

UNSHARP says that having unsharp degrees of belief is compatible with perfect rationality. If UNSHARP were true, there would be a good answer to the question:

> How do unsharp probabilities constrain rational action?

But there is no good answer to that question. Permissive rules are too permissive: they wrongly say that it can be rational to reject both bets in the Bet A/Bet B situation. Strict rules are too strict: in pinning down precise betting odds, they undercut a main motivation for introducing unsharp probabilities in the first place. And global rules such as NARROW, PLAN, and SEQUENCE fail as well.

So UNSHARP is false. Perfect rationality requires perfectly sharp probabilities.[27]

27. Thanks to Brian Weatherson, Cian Dorr, Andy Egan, Tom Dougherty, Susanna Rinard, Bas van Fraassen, James Joyce, Jim Pryor, Teddy Seidenfeld, Alan Hájek, Boris Kment, Paul Weirich, an anonymous *Philosophers' Imprint* referee, participants in the 2007 "MITing of the Minds" conference, a reading group at the University of British Columbia, the Corridor Group, and participants at a graduate seminar meeting at the University of Michigan. Special thanks to Caspar Hare for his incisive conference comments, and to David Christensen for a prompt and helpful response to a last-minute philosophical SOS.

References

Michael Bratman. Planning and the stability of intention. *Minds and Machines*, 2:1–16, 1992.

David Christensen. Clever bookies and coherent beliefs. *Philosophical Review*, pages 229–247, 1991.

Laura DeHelian and Edward F. McClennen. Planning and the stability of intention: a comment. *Minds and Machines*, 3:319–333, 1993.

Tom Dougherty, February 2007. Personal communication.

Richard Feldman. Reasonable religious disagreements. In Louise M. Antony, editor, *Philosophers without Gods: Meditations on Atheism and the Secular Life*. Oxford University Press, 2007.

Peter Gärdenfors and Nils-Eric Sahlin. Unreliable probabilities, risk taking, and decision making. *Synthese*, 53:361–386, 1982.

David Gauthier. *Morals by Agreement*. Oxford Clarendon Press, 1986.

David Gauthier. Resolute choice and rational deliberation: A critique and a defense. *Nous*, 31(1):1–25, 1997.

David Gauthier. Rethinking the Toxin Puzzle. In Jules L. Coleman, editor, *Rational Commitment and Social Justice: Essays for Gregory Kavka*. Cambridge University Press, 1998.

Itzhak Gilboa and David Schmeidler. Maximin expected utility theory with non-unique prior. *Journal of mathematical economics*, 18:141–153, 1989.

I. J. Good. Rational decisions. *Journal of the Royal Statistical Society, Series B*, 14(1):107–114, 1952.

Alan Hájek. Scotching Dutch Books? *Philosophical Perspectives*, 19:139–151, 2005.

Alan Hájek. Dutch Book arguments. In Paul Anand, Prasanta Pattanaik, and Clemens Puppe, editors, *Oxford Handbook of Rational and Social Choice*, pages 173–195. Oxford University Press, 2008.

Caspar Hare. What the mushy credences person should say: comments on "subjective probabilities should be sharp". Unpublished comments delivered at MITing of the Minds conference, January 2007.

Richard C. Jeffrey. Bayesianism with a human face. In John Earman, editor, *Testing scientific theories*, volume 10 of *Minnesota Studies in the Philosophy of Science*, pages 133–156. University of Minnesota Press, Minneapolis, 1983.

James M. Joyce. How probabilities reflect evidence. *Philosophical Perspectives*, 19:153–178, 2005.

Mark Kaplan. *Decision Theory as Philosophy*. Cambridge University Press, Cambridge, 1996.

Mark Kaplan. In defense of modest probabilism. *Synthese*, March 2009. doi: 10.1007/s11229-009-9483-2.

John G. Kemeny. Fair bets and inductive probabilities. *Journal of Symbolic Logic*, 20(3):263–273, 1955.

John Maynard Keynes. The general theory of employment. In *The general theory and after: part II defence and development*, volume 14 of *The collected writings of John Maynard Keynes*, pages 109–124. Macmillan, London, 1973.

Henry Kyburg. Subjective probability: Criticisms, reflections, and problems. *Journal of Philosophical Logic*, 7:157–180, 1978.

Isaac Levi. *The enterprise of knowledge*. MIT Press, Cambridge, 1980.

Isaac Levi. Imprecision and indeterminacy in probability judgment. *Philosophy of Science*, 52(3):390–409, 1985.

Isaac Levi. The demons of decision. *The Monist*, 70:193–211, 1987.

Edward F. McClennen. *Rationality and dynamic choice: foundational explorations*. Cambridge University Press, 1990.

Edward F. McClennen. The rationality of being guided by rules. In Alfred R. Mele and Piers Rawling, editors, *Oxford Handbook of Rationality*, Oxford Handbooks in Philosophy, pages 222–239. Oxford University Press, 2004.

Mark J. Schervish, Teddy Seidenfeld, and Joseph B. Kadane. Two measures of incoherence: how not to gamble if you must. Preprint available at PhilSci Archive, 1998. URL `http://philsci-archive.pitt.edu/archive/00000160/`.

Teddy Seidenfeld. Review of Paul Horwich's "Probability and Evidence". *Philosophical Review*, 93:474–483, 1984.

Teddy Seidenfeld. When normal and extensive form decisions differ. In D. Prawitz, B. Skyrms, and D. Westerahl, editors, *Logic, methodology and philosophy of science*, volume 9, pages 451–463. Elsevier Science, 1994.

Teddy Seidenfeld, Mark J. Schervish, and Joseph B. Kadane. Decisions without ordering. In W. Sieg, editor, *Acting and reflecting: the interdisciplinary turn in philosophy*, pages 143–170. Kluwer Academic Publishers, 1990.

J. J. C. Smart. Extreme and restricted utilitarianism. *Philosophical Quarterly*, 6(25):344–354, 1956.

Scott Sturgeon. Reason and the grain of belief. *Noûs*, 42:139–165, 2008.

Bas C. van Fraassen. Vague expectation value loss. *Philosophical Studies*, 127:483–491, 2006.

Peter Walley. *Statistical reasoning with imprecise probabilities*. Chapman and Hall, 1991.

Brian Weatherson. Should we respond to evil with indifference? *Philosophy and Phenomenological Research*, 70:613–635, 2005.

Brian Weatherson. Decision making with imprecise probabilities. Manuscript, November 2008. URL `http://brian.weatherson.org/vdt.pdf`.

Paul Weirich. Risk's place in decision rules. *Synthese*, 126:427–441, 2001.

Roger White. Epistemic permissiveness. *Philosophical Perspectives*, 19 (1):445–459, 2005.

Marco Wolfenson and Terrence L. Fine. Bayes-like decision making with upper and lower probabilities. *Journal of the American Statistical Association*, 77(377):80–88, 1982.