On choosing how to choose*

Richard Pettigrew

February 5, 2023

Abstract

A decision theory is self-recommending if, when you ask it which decision theory you should use, it considers itself to be among the permissible options. I show that many alternatives to expected utility theory are not self-recommending, and I argue that this tells against them.

Contents

*This is the second draft of this paper; the first can still be found at PhilArchive here. A Mathematica notebook containing all the calculations carried out in this paper is available here. I'm very grateful to Liam Kofi Bright for extremely helpful comments—they greatly improved the paper.

It is our lot to face decisions when we have no certainty about which option from among those available to us will lead to the best outcome. Uncertain about the day's weather, you must choose whether or not to wear a raincoat when you leave your house; uncertain about what it would help them most to hear, I must choose what to say to a friend who is going through a bad time; and uncertain what effect different approaches will have, a parent must choose how to raise their child. How are we to make such choices? It is the task of decision theory to provide an answer. And philosophers, economists, and psychologists have met this remit by developing a slew of rival theories of rational decision. Expected utility theory is the most well known and widely used, but there are many alternatives available, and we will meet a good few in the course of this paper.

Which of these theories should you use? You might argue for a particular decision theory directly in a number of ways. You might note that it agrees with your intuitive verdict about a specific decision problem that you describe, while its rivals don't; or you might note that your favoured theory has a formal feature that you find intuitively desirable, while its rivals lack that feature. But there is another natural approach, and it has the advantage that it avoids such appeals to our intuitive judgments. It begins with the observation that a decision theory is an account of rational means-ends reasoning: agnostic about whether your ends are good or bad, desirable or undesirable, benevolent or malign, it purports to tell you the rational way to pick between different possible means to the ends you in fact have. Granted this, it seems natural to assess a decision theory by asking how well it performs in the role of getting you those ends. The only problem with this approach is that, in order to assess a decision theory or anything else as a means to your ends, we need an account of which means to your ends it is rational to use. And without a decision theory, we don't have that.

Yet all is not lost, for this line of thinking nonetheless furnishes us with a test we can conduct on a theory of decision-making, and while it might not tell in favour of the theory if it passes, it seems to tell against it if it fails. We can ask of the theory: If I were to use you not only to make my normal day-to-day decisions, but also to make the higher-order decision about which decision theory to use, would you recommend yourself? If it would, we call it *self-recommending*; if it wouldn't, we call it *self-undermining*. I claim that no self-undermining decision theory can be correct. This is not to say that a self-recommending theory is thereby adequate—the theory that says that any available option is rationally permissible is self-recommending, but it is not correct. Nonetheless, we can use this test to winnow the list of candidate decision theories, removing those that fail it.

Below, I describe a number of putative decision theories and conduct this test on them. As we will see, many of the most popular rivals to expected utility theory are not self-recommending. In Section 2, we will see that Savage's formulation of expected utility theory is self-recommending; Section 3 shows that Lara Buchak's risk-weighted expected utility theory is self-undermining, while Section 4 shows that nearly all the putative decision theories for imprecise credence are self-undermining; in Section 5, we will see that many decision theories designed for use under Knightian uncertainty are self-recommending. In **??**, I motivate the test further, and I respond to certain objections to my claim that an adequate decision theory must be self-recommending.

1 Why is it bad to be self-undermining?

Before we meet those results, let me say more about why a decision theory should be self-recommending. And to do that, let me begin by expanding on the ways I enumerated above in which we might adjudicate between such theories. There are three: (i) adequacy to specific intuitions; (ii) adequacy to general intuitions; and (iii) by-their-fruits considerations. (If you are already convinced that self-undermining decision theories should be rejected, you can skip to Section 2 without loss.)

On the first, we describe a decision problem, elicit our intuitive reaction to it, and find in favour of decision theories that agree with that reaction and against those that disagree with it. When we appeal to the Allais preferences to motivate a risk-sensitive alternative to expected utility theory, we do this (Allais, 1953; Buchak, 2022). Intuitively, it is rationally permissible to have the Allais preferences; expected utility theory says it isn't, but risk-weighted expected utility theory says it is.

On the second, we describe a general feature that a decision theory might have, elicit our intuitive reaction to it, and find in favour of decision theories that boast it and against those that don't. When we note that risk-weighted expected utility theory does not satisfy the so-called Independence Axiom of decision theory, but expected utility theory does, and argue in favour of the latter on that basis, we do this (von Neumann & Morgenstern, 1947).

Finally, the third way to adjudicate between two decision theories. One problem with the previous two ways, both of which appeal to our intuitions, is that they are apt to lead us to a stalemate. For you, perhaps, the intuition that the Allais preferences are rational is stronger than your intuition that the Independence Axiom is true; for me, perhaps, it's the other way around. And indeed the tension might be internal: perhaps you feel the intuitive pull of both, but you can't determine which is stronger.

How are we to resolve such disagreements? We need some uncontested ground on which we might stand to arbitrate. We might hope that all parties to the dispute agree that we adopt a decision theory in order to get as much of what we want as we can. And so we might hope to agree to settle the dispute between our decision theories by looking at how well they serve that purpose. That is, we might hope to assess our principles of means-end reasoning—for that is all a decision theory really is—by asking how well they perform as means to our ends. As the gospel of Matthew has it, ye shall know them by their fruits.

The classic examples of this approach are what I will call exploitability arguments, which include money pump and diachronic incoherence arguments (Gustafsson, 2022). Consider Maximin, which says that you should pick from among the options whose worst-case outcome is best (Wald, 1945). It is often pointed out that this decision theory is diachronically incoherent. That is, there are certain pairs of decisions you might face such that, if you were to choose using Maximin in both cases, then you would end up picking a particular pair of options, one from each decision problem, when there was an alternative pair, one from each decision problem, that would result in greater total utility for sure—in the jargon of decision theory, the alternative pair of options strongly dominates the pair that Maximin requires you to choose; that is, the alternative pair is strictly better no matter what. For instance, suppose you first face a decision problem with two options: the first option gives utility 1 for sure, while the second gives utility 3 if it will rain tomorrow and utility 0 otherwise. Maximin will tell you to choose the sure thing. Next, suppose you face another decision problem with two options: again, the first gives utility 1 for sure, while the second gives utility 0 if it will rain tomorrow and utility 3 otherwise. Maximin again demands the sure thing. However, if you had instead chosen the second option in both cases, rather than the sure thing, you would have ended up with utility $3 + 0 = 3$ if it were to rain tomorrow, and $0 + 3 = 3$ otherwise. So, the argument goes, Maximin is a bad theory of means-end rationality because, in certain situations, it leads you to make choices that don't get you as close to your ends as other choices would have. And note: this doesn't beg the question against Maximin, since that decision theory deems strongly dominated options impermissible.

Other examples are easy to come by:

- *The Money Pump Argument for Non-Cyclical Preferences* (Gustafsson, 2022) Suppose you strictly prefer B to A, C to B, and A to C. So you have cyclical preferences, $A \prec B \prec C \prec A$. Then I begin by offering you a choice between A and B. In line with your preferences, you choose B. Then I offer you a choice between sticking with B or paying a small amount to swap to C. In line with your preferences, you choose to swap. Then, finally, I offer you a choice between sticking with C or paying a small amount to swap to A. In line with your preferences, you choose to swap. You are then left with A less the money you paid for the two swaps, when you could have ended up with A simply by choosing that when offered the initial choice between A and B. So there is an alternative sequence of choices you might have made that dominates the one you did make.

- *The Diachronic Incoherence Argument for Exponential Discounting* (Strotz, 1955) Suppose you discount the future using a discount function that isn't exponential. Then there will be a pair of options A and B such that you strictly prefer A to B at one time, but then prefer B to A at a later time. So I first offer you a choice between A or B at the earlier time. In line with your preferences at that time, you choose A. Then, at the later time, I offer you the option to swap from B to A for a small price. In line with your preferences, you choose to swap. Again, if you'd simply chosen A at the earlier time, you would have been better off for sure.

I want to persuade you there is a problem with this sort of argument. What's more, if we try to fix it, it leads to the sort of argument I will present in this paper. Take the diachronic incoherence argument against Maximin, for instance. It identifies a very specific sort of situation in which that decision theory leads to an outcome that is bad even by its own lights. It specifies a particular pair of decision problems you might face in sequence, and it shows that, if you do, Maximin will lead you to choose a strongly dominated sequence of options. But the argument speaks only of this specific sequence of decision problems, or at most others structurally like it. It says nothing about how well Maximin performs in other situations, when faced with other sequences of decision problems. It does not rule out the possibility that, whatever vices the decision theory displays when faced with that pair of decision problems, or others structurally like it, it displays virtues when faced with other decision problems, and those virtues outweigh the vices. Poor performance in a particular situation is not usually reason to abandon a tool, since it might compensate by performing well in other situations. Yet that is exactly what the diachronic coherence argument—as well as other exploitability arguments—asks us to do.

So, if we wish to evaluate decision theories by their performance when faced with decision problems, we have to make precise exactly how we weigh performance in some situations against performance in other situations to determine whether one compensates for the other. How do we do that? Well, that is exactly the sort of question a decision theory is supposed to answer. It takes in the performance of different options in different situations and weighs those against each other to determine which are optimal. So, to ensure that we are not begging the question against a decision theory, we must use it—the very decision theory we are judging—to judge whether it is the best tool we can use when we are trying to optimise our utility. That is, we must ask whether it is self-recommending.

That's one way we might motivate the test for decision theories that I will explore in this paper. Here's another. It's just a straightforward argument that self-undermining decision theories cannot be correct. A decision theory categorises available options into those that are rationally permissible for a particular agent and those that are rationally impermissible for that agent. Now suppose the decision theory is correct: that is, the options it declares permissible are indeed exactly those that are permissible for the agent. But one of its declarations is that using it is itself impermissible—that's what makes it self-undermining. Now, how could it ever be rationally impermissible to use a decision theory that gives the correct verdicts about rational permissibility? Surely it couldn't! And if that's so, then the decision theory cannot be correct, and we have a contradiction.

Now, perhaps you might reply that it could be impermissible to use a decision theory that gives the correct verdicts if that decision theory were so complex and computationally demanding that it lay beyond the agent's cognitive capacities—not much good using something that gives the right verdicts every time if you can never discover what verdicts it gives. But that is not the case we are considering here. Indeed, we will assume explicitly that, if I use a decision theory when I face a particular decision problem, I will pick one of the options it deems rationally permissible.

So, in conclusion: it can never be irrational for an ideal agent to use the correct account of rationality to make their decisions, and so a self-undermining decision theory cannot give the correct account of rationality.

Before we move on, let me point to another advantage of my approach over the exploitability arguments. Such an argument always describes a *sequence* of decision problems and shows that, when you use the decision theory to face each problem in that sequence, you will choose a sequence of options that is dominated by another—that is, you will end up *diachronically* incoherent. But some philosophers don't think diachronic incoherence is a marker of irrationality. For instance, they think that each of the decisions in the sequence is made by a different self or timeslice, and a collection of selves or timeslices losing money for sure is no indication of the irrationality of any of those individuals, and perhaps not even an indication that the agent of which they are parts is irrational.[1] But I needn't worry about that. My arguments show that, even when there is only a *single* decision problem in play, certain decision theories fail to recommend themselves.

[1]David Christensen (1991) argues that the standard diachronic betting argument in favour of Bayesian conditionalization fails for this reason.

2 Expected utility theory

2.1 What is decision theory?

What is a decision theory? Informally, it is a theory that tells you what to choose when faced with a decision problem. What, then, is a decision problem? Informally, it's a collection of options that are available for you to choose. More formally, for our purposes here, you begin with a set of possible states of the world. These states are jointly exhaustive and pairwise exclusive, so that the world must be exactly one of these ways—as it is sometimes said, the states of the world partition the logical space. Each option in the decision problem is defined by giving its utility at each state of the world, which measures how much you value being in that state of the world having chosen that option. So if $\mathcal{S}$ is the set of states of the world, $\mathcal{O}$ is the set of options defined on $\mathcal{S}$ that are available in this decision problem, and if o is in $\mathcal{O}$ and s is in $\mathcal{S}$, then $o(s)$ is a real number that gives the utility of picking option o when the state of the world is s. So, in the decision whether to take your raincoat, the states might be *Rain* and *Not Rain*, and the options might be *Take the raincoat* and *Leave the raincoat*, and the utility of taking the raincoat at the state where it rains might be 3, so that $Take(Rain) = 3$, while leaving the raincoat at that state might be -1, so that $Leave(Rain) = -1$.

Having said what a decision problem is, let me return to the question of what a decision theory is. For our purposes here, it is a function. As input, it takes a decision problem and possibly some of your attitudes, such as your credences, if you have them, or your attitudes to risk, if you have those. As output, it returns a subset of the options available in the decision problem. This subset is the set of options that it deems rationally permissible for someone who faces that decision problem and has those attitudes.

2.2 What is expected utility theory?

Let's see this at work in the case of standard expected utility theory.[2] I'll introduce a simple example here and deploy it throughout the paper to illustrate different decision theories. Suppose that I have been experiencing various medical symptoms, and my doctor is sure I have one of two illnesses, but they're unsure which one. Two treatments are available. If I take Medicine 1 and it's Illness 1, I'll be almost completely cured, but if it's Illness 2, the treatment will have no effect. If I take Medicine 2, and it's Illness 1, I'll become worse, but if it's Illness 2, I'll be completely cured. The following table shows my utilities for each option at each state of the world, where 3 is my current level of utility:

	i_1	i_2
m_1	9	3
m_2	0	10

So each row corresponds to an option—the two medicines, m_1 and m_2—and each column corresponds to a state of the world—the two possible illnesses, i_1 and i_2. In what follows, we'll often denote a decision with these payouts by the quadruple $(9, 3; 0, 10)$. After listening to the doctor's reasoning, I'm 40% confident I have Illness 1 and 60% confident I have Ilness

[2] For the cognoscenti, the version I'll present here is of the sort described by Savage (1954), in which it is assumed that the options are independent of the states of the world. This assumption is dropped in the evidential decision theory of Jeffrey (1965) and the causal decision theory of Stalnaker (1972 [1981]); Gibbard & Harper (1978); Joyce (1999).

2. According to expected utility, I should decide which treatment to undertake by maximising expected utility from the point of view of my credence function, which I denote P. The expected utility of m_1 is

$$\text{Exp}_P(m_1) = P(i_1)m_1(i_1) + P(i_2)m_1(i_2) = (0.4 \times 9) + (0.6 \times 3) = 5.4$$

while the expected utility of T_2 is

$$\text{Exp}_P(m_2) = P(i_1)m_2(i_1) + P(i_2)m_2(i_2) = (0.4 \times 0) + (0.6 \times 10) = 6$$

So, according to expected utility theory, I should begin Treatment 2.

In general, suppose $\mathcal{S} = (s_1, \ldots, s_n)$ are the states of the world, P is a credence function that assigns a probability to each of those states, and o is an option defined over those states.[3] Then the expected utility of o from the point of view of P is:

$$\text{Exp}_P(o) = P(s_1)o(s_1) + \ldots + P(s_n)o(s_n)$$

That is, it is the sum of the probability-weighted utilities of the option at the different states of the world. We can then state the decision theory:

> **Expected Utility Theory$_P$ (EU$_P$)** Given an option o in $\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that $\text{Exp}_P(O) < \text{Exp}_P(O')$.

2.3 Is expected utility theory self-recommending?

Now, suppose we wish to ask expected utility theory whether it would recommend itself. To do this, we must describe a decision problem in which expected utility theory and its rivals are the options, and the states of the world specify which decision problem you will face using your chosen decision theory, along with enough information to determine how much utility you'll obtain from using a particular decision theory when faced with that decision problem. I'll call this our *meta decision problem*, and I'll call the states of the world on which the options are defined the *meta states of the world*. Given a particular decision theory, and given whichever of your attitudes it needs to give a verdict, the utility of using it with those attitudes at one of these meta states of the world is the utility it obtains for you at that state of the world; that is, it is the utility, at that state of the world, of whichever option from the decision problem the decision theory tells you to choose given your attitudes.

To keep things simple, we'll assume you know you'll face a decision problem in which there are exactly two states of the world, s_1 and s_2, and exactly two options, o_1 and o_2. That is, you'll face a decision problem with two options, o_1 and o_2, each defined on two states of the world, s_1 and s_2. We'll write this $(a, b; c, d)$, where a is the utility of o_1 at s_1, b is the utility of o_1 at s_2, c is the utility of o_2 at s_1, and d is the utility of o_2 at s_2. Like Savage, we'll assume that utilities are bounded above and below, and we'll measure utility on a scale on which 0 is the lowest possible utility and 1 is the highest. So $0 \leq a, b, c, d \leq 1$. Each meta state of the world in our meta decision problem will specify which of these decision problems you'll face.

In order to determine what utility a decision theory has when faced with such a decision problem, we clearly need to specify which of the original states of the world we're in—i.e. s_1

[3] Throughout, we assume that P is a probability function, so that $P(s_1) + \ldots + P(s_n) = 1$.

or s_2. But in fact this is not enough. After all, a decision theory only tells you which options are permissible, and there might be more than one. For instance, in the binary decisions on which we're focusing here, if o_1 and o_2 have the same expected utility from the point of view of P, then EU_P says both are permissible. So our meta states of the world must specify not only the decision problem you'll face and the state of the world you're in, but also the option you'll in fact choose if you may choose either. Let t_1 be the proposition that, in the case of a tie, you'll choose option o_1, and let t_2 be the proposition that, in the case of a tie, you'll choose o_2.

So a meta state of the world has the following form:

$$\langle (a, b; c, d), s_i, t_j \rangle$$

It is a triple consisting of a decision problem $(a, b; c, d)$, a state of the world s_i, and a tie-breaker result t_j. And the utility of EU_P at such a state of the world is:

- a, if *either* $\mathrm{Exp}_P(o_2) < \mathrm{Exp}_P(o_1)$ and $i = 1$, *or* $\mathrm{Exp}_P(o_2) = \mathrm{Exp}_P(o_1)$, $i = 1$, and $j = 1$;
- b, if *either* $\mathrm{Exp}_P(o_2) < \mathrm{Exp}_P(o_1)$ and $i = 2$ *or* $\mathrm{Exp}_P(o_2) = \mathrm{Exp}_P(o_1)$, $i = 2$, and $j = 1$.
- c, if *either* $\mathrm{Exp}_P(o_1) < \mathrm{Exp}_P(o_2)$ and $i = 1$ *or* $\mathrm{Exp}_P(o_2) = \mathrm{Exp}_P(o_1)$, $i = 1$, and $j = 2$.
- d, if *either* $\mathrm{Exp}_P(o_1) < \mathrm{Exp}_P(o_2)$ and $i = 2$ *or* $\mathrm{Exp}_P(o_2) = \mathrm{Exp}_P(o_1)$, $i = 2$, and $j = 2$.

In order to ask EU_P whether it recommends itself, we need to extend the probability function P, which assigns credences only to the states of the world, to a credence function $P^\star$, which assigns credences to the meta states of the world. We'll assume that, according to $P^\star$, which decision problem you'll face is independent of the original states of the world and both are independent of the tie-breaker proposition. Your credences in the original states remain the same: $P^\star(s_1) = P(s_1) = p$ and $P^\star(s_2) = P(s_2) = 1 - p$. However you'll in fact break a tie should you face one, your current credences that you'll break them one way or the another are $P^\star(t_1) = t$ and $P^\star(t_2) = 1 - t$. And we'll assume you pick a measure μ over the possible decision problems D of the form $(a, b; c, d)$—we'll let $\mathcal{D}$ be the set of all such decision problems. We write $P^\star_{\mu,t}$ for this probabilistic credence function. So, if $\mathcal{D}'$ is a measurable subset of $\mathcal{D}$, then, for instance,

$$P^\star_{\mu,t}(D \in \mathcal{D}' \& \, s_1 \,\&\, t_1) = t \times p \times \mu(D \in \mathcal{D}')$$

Then, given a decision theory DT, the expected utility of DT from the point of view of P and μ is:

$$\mathrm{Exp}_{P^\star_{\mu,t}}(\mathrm{DT}) = \sum_{i=1}^{2} \sum_{j=1}^{2} P(s_i) P^\star_{\mu,t}(t_j) \int_{\mathcal{D}} \mathrm{DT}(D, t_j)(s_i) \, d\mu$$

where $\mathrm{DT}(D, t_j)$ is the option that DT would lead you to choose if faced with decision problem D and if the tie-breaker proposition were t_j, and $\mathrm{DT}(D, t_j)(s_i)$ is the utility of this option at state s_i. Then it's reasonably easy to see that EU_P is self-recommending for any μ. After all,

$$\mathrm{Exp}_{P^\star_{\mu,t}}(\mathrm{EU}_P(D, t_j)) \geq \mathrm{Exp}_{P^\star_{\mu,t}}(\mathrm{DT}(D, t_j))$$

for any decision theory DT. So

$$\begin{aligned}
\mathrm{Exp}_{P^\mu}(\mathrm{EU}_P) &= \sum_{i=1}^{2}\sum_{j=1}^{2} P(s_i)P^\star_{\mu,t}(t_j)\int_{\mathcal{D}} EU_P(D,t_j)(s_i)\,d\mu \\
&= \sum_{j=1}^{2} P^\star_{\mu,t}(t_j)\int_{\mathcal{D}} \sum_{i=1}^{2} P(s_i)\mathrm{EU}_P(D,t_j)(s_i)\,d\mu \\
&= \sum_{j=1}^{2} P^\star_{\mu,t}(t_j)\int_{\mathcal{D}} \mathrm{Exp}_P(\mathrm{EU}_P,t_j)\,d\mu \\
&\geq \sum_{j=1}^{2} P^\star_{\mu,t}(t_j)\int_{\mathcal{D}} \mathrm{Exp}_P(\mathrm{DT},t_j)\,d\mu \\
&= \sum_{j=1}^{2} P^\star_{\mu,t}(t_j)\int_{\mathcal{D}} \sum_{i=1}^{2} P(s_i)\mathrm{DT}(D,t_j)(s_i)\,d\mu \\
&= \sum_{i=1}^{2}\sum_{j=1}^{2} P(s_i)P^\star_{\mu,t}(t_j)\int_{\mathcal{D}} \mathrm{DT}(D,t_j)(s_i)\,d\mu \\
&= \mathrm{Exp}_{P^\mu}(DT)
\end{aligned}$$

So:

Theorem 1. *For any P, μ, and t,* EU_P *is permissible by the lights of* $\mathrm{EU}_{P^\star_{\mu,t}}$.

So expected utility theory is self-recommending.

Indeed, in Figure 1, we plot, for $p = 0.1, 0.2, \ldots, 0.8, 0.9$, the expected utility, from the point of view of a probability function that assigns p to s_1 and $1-p$ to s_2, of using expected utility theory with a probability function that assigns q to s_1 and $1-q$ to s_2 as your decision rule, when you know you'll face a decision problem of the form $(a, b; c, d)$ for $0 \leq a, b, c, d \leq 1$, and you place a uniform distribution over these possible decision problems. As you can see from the figure, the expected utility from the point of view of p of using expected utility theory with q is maximal for $p = q$.

It is worth noting that this reasoning only holds for Savage's (1954) version of expected utility theory, in which the states of the world are independent of the options chosen. It is a further question whether the causal or evidential versions of expected utility theory are self-undermining in the current sense. I leave that for future work.

3 Decision theories that are sensitive to risk

For individuals whose beliefs are represented by precise credences, expected utility theory is not the only normative decision theory that has been proposed (Buchak, 2022). In response to the so-called Allais paradox, decision theorists have presented a range of alternatives that permit individuals to combine their credences with their attitudes to risk in certain ways in order to make decisions (Kahneman & Tversky, 1979; Machina, 1982; Quiggin, 1982). In this section, I'll consider Lara Buchak's *risk-weighted expected utility theory* (Buchak, 2013).

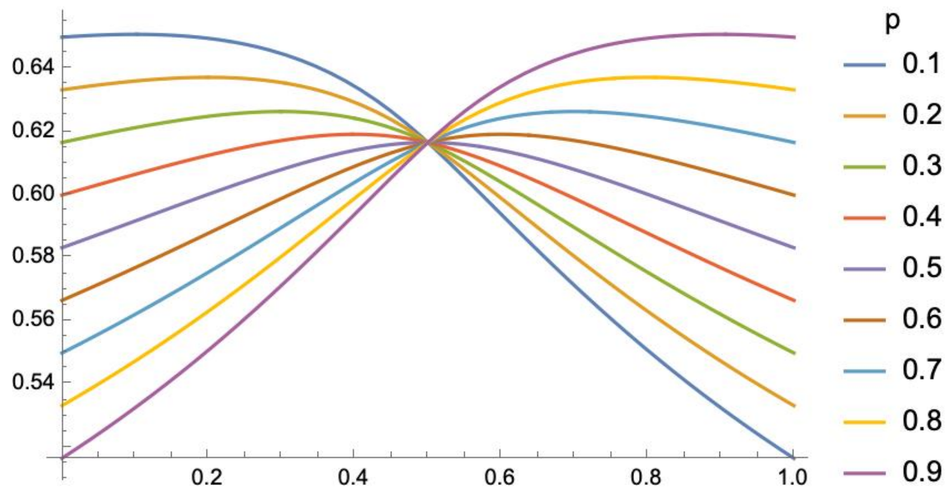

Figure 1: For $p = 0.1, 0.2, \ldots, 0.8, 0.9$, we plot $\text{Exp}_p(\text{EU}_q)$, for $0 \leq q \leq 1$.

3.1 What is risk-weighted expected utility theory?

It's difficult to say precisely what it means to be risk-averse or risk-inclined or risk-neutral without assuming a particular theory of the phenomenon, but speaking roughly we might say this: you are risk-averse if, in the evaluation of an option, you give greater weight to the utilities it will obtain for you in its worst-case outcomes and less weight to the utilities it will obtain for you in its best-case outcomes than expected utility theory tells you to give; you're risk-inclined if you give greater weight to the best-case outcomes and less to the worst-case than expected utility theory requires; and you're risk-neutral if you give exactly the weights that expected utility theory suggests.

So, for instance, suppose that your utility is linear in some commodity—money or chocolate or the number of happy years of life lived, for instance. That is, acquiring a particular quantity of that commodity—a dollar, a chocolate bar, a happy year of life, perhaps—increases your utility by the same amount regardless of how much of that commodity you already have. And now suppose I offer you 1 unit of that commodity for sure, or a fair coin toss that gives you 3 units if the coin lands heads or 0 units if it lands tails. Then expected utility theory will require you to take the coin toss, since its expected utility is greater than the sure thing; but a risk-averse individual might take the sure thing because, in the evaluation of the coin toss, they don't give equal weight to both outcomes, as expected utility theory tells you to do, and instead give greater weight to the worst-case outcome in which you receive 0 units and less to the best-case outcome in which you receive 3 units.

To see how Lara Buchak's risk-weighted expected utility theory (henceforth, REU) tells you to weight the different outcomes of the different options available in a decision problem in order to choose between them, let's first see how expected utility theory (EUT) does it. EUT tells you that the value of an option o, defined on finitely many states of the world, $s_1, \ldots, s_n$, is its expected utility:

$$\text{Exp}_P(o) = P(s_1)o(s_1) + \ldots + P(s_n)o(s_n)$$

But note that, if we order the states of the world by their utilities, so that $o(s_{i_1}) \leq \ldots \leq o(s_{i_n})$, then we can also write this as follows:

$$
\begin{aligned}
\mathrm{Exp}_P(o) = P(s_{i_1} \vee \ldots \vee s_{i_n})o(s_{i_1}) + \\
P(s_{i_2} \vee \ldots \vee s_{i_n})[o(s_{i_2}) - o(s_{i_1})] + \\
P(s_{i_3} \vee \ldots \vee s_{i_n})[o(s_{i_3}) - o(s_{i_1})] + \ldots \\
\ldots + \\
P(s_{i_{n-1}} \vee s_{i_n})[o(s_{i_{n-1}}) - o(s_{i_{n-2}})] + \\
P(s_{i_n})[o(s_{i_n}) - o(s_{i_{n-1}})]
\end{aligned}
$$

That is, the expected utility of o is obtained as follows: first, take the smallest possible utility that o will obtain for you—that is, $o(s_{i_1})$—and weight it by the probability that o will obtain at least that for you—that is, $P(s_{i_1} \vee \ldots \vee s_{i_n})$, which is obviously 1; next, add to this the difference between the smallest possible utility and the next smallest possible utility—that is, $o(s_{i_2})$—and weight that difference by the probability that o will obtain at least that next smallest possible utility—that is, $P(s_{i_2} \vee \ldots \vee s_{i_n})$; and then carry on in this way until you reach the greatest possible utility that o will obtain for you.

The risk-weighted expected utility of o is then given in almost exactly the same way, except that the weights are given not by the probabilities, but by the probabilities skewed by a risk function R that represents your attitudes to risk. R takes probabilities in $[0, 1]$ and returns new weights in $[0, 1]$. We assume R is a strictly increasing, continuous function with $R(0) = 0$ and $R(1) = 1$. Then the risk-weighted expected utility of option o is given as follows:

$$
\begin{aligned}
\mathrm{RExp}_P(o) = R(P(s_{i_1} \vee \ldots \vee s_{i_n}))o(s_{i_1}) + \\
R(P(s_{i_2} \vee \ldots \vee s_{i_n}))[o(s_{i_2}) - o(s_{i_1})] + \\
R(P(s_{i_3} \vee \ldots \vee s_{i_n}))[o(s_{i_3}) - o(s_{i_1})] + \ldots \\
\ldots + \\
R(P(s_{i_{n-1}} \vee s_{i_n}))[o(s_{i_{n-1}}) - o(s_{i_{n-2}})] + \\
R(P(s_{i_n}))[o(s_{i_n}) - o(s_{i_{n-1}})]
\end{aligned}
$$

Risk-Weighted Expected Utility Theory$_{R,P}$ ($REU_{R,P}$) Given an option o in $\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that $\mathrm{RExp}_{R,P}(o) < \mathrm{RExp}_{R,P}(o')$.

To illustrate, return to the example of the sure thing versus the coin toss from above. Suppose the utility you assign to having n units of the commodity in question is n. So, using the notation from the previous section, we can describe the decision problem like this: $(1, 1; 0, 3)$. And suppose $P(s_1) = P(s_2) = 0.5$, and $R(p) = p^2$. Then the risk-weighted expected utility of option o_1 is:

$$\mathrm{RExp}_{P,R}(o_1) = 1 + R(0.5)[1 - 1] = 1 + (0.25 \times 0) = 1$$

And for option o_2 it is:

$$\mathrm{RExp}_{P,R}(o_2) = 0 + R(0.5)[3 - 0] = 0 + (0.25 \times 3) = 0.75$$

So if your risk function is $R(p) = p^2$, then you should take the sure thing in that case.[4]

In general, if R is convex, then it gives rise to risk-averse behaviour; if R is concave, it gives risk-inclined behaviour; and if R is simply the identity function, so that $R(p) = p$, then the risk-weighted expected utility of an option is simply its expected utility, and so REU agrees with EU, and we say that it gives rise to risk-neutral behaviour.

3.2 Is risk-weighted expected utility theory self-recommending?

So much for what the theory says. Now we must ask what it thinks of its own performance. We assume that our decision-maker knows that they will face a binary decision problem of the form: $(a, b; c, d)$, where $0 \leq a, b, c, d, \leq 1$. Given an instance of Buchak's theory, determined by a credence function P and a risk function R, does that version recommend itself? That is, given the meta decision problem in which the meta states of the world have the form, $\langle (a, b; c, d), s_i, t_j \rangle$, and in which the options are all the available decision theories, and given a version of Buchak's theory, which decision theory or theories does it deem permissible, and does it include itself among them? I'm afraid the answer to the latter question is that it doesn't.

The version of Buchak's theory that I laid out above, which is the version that she presents herself, is formulated for decision problems in which the options are defined over only finitely many states of the world. But there are clearly continuum-many different *meta* states of the world. So, we need a version that is formulated for such decision problems. Fortunately, the natural generalization of Buchak's theory can be given by analogy with the generalization that Quiggin gives of his own rank-dependent utility theory (Quiggin, 1982).[5]

Suppose that the utilities offered by a particular option o are bounded above by U_1 and below by U_0. And, for $U_0 \leq u \leq U_1$, let $o > u$ be the proposition that says that the utility of o is greater than u. That is, it is the proposition that is true at exactly those states s such that $o(s) > u$. Then the expected utility of o from the point of view of P is equal to the following:

$$\text{Exp}_P(o) = U_0 + \int_{u=U_0}^{U_1} P(o > u)\, du$$

And the risk-weighted expected utility of o is defined as follows:

$$\text{RExp}_{R,P}(o) = U_0 + \int_{u=U_0}^{U_1} R(P(o > u))\, du$$

As before, we'll assume that the three components of the meta state—the decision problem you'll face, the original state of the world, and the tie-breaker proposition—are all independent. Then, we fix your probability function P over s_1 and s_2, which assigns $P(s_1) = p$ and $P(s_2) = 1 - p$, and your measure μ over the possible values of a, b, c, d, and we combine

[4]Note that, for this reason, REU with $R(p) = p^2$ is diachronically incoherent. In Section 1, we saw that, faced with $(1, 1; 0, 3)$ and $(1, 1; 3, 0)$, Maximin requires you to choose $(1, 1)$ then $(1, 1)$, but $(0, 3)$ then $(3, 0)$ strongly dominates it. We can see now that, if $P(s_1) = P(s_2) = 0.5$, then REU with P and $R(p) = p^2$ requires the same. And indeed, for all risk functions R, if $R(p) \neq p$ for some p, then there is P such that REU with P and R is diachronically incoherent.

[5]Thanks to Kenny Easwaran for pointing me to this. Quiggin's theory stands to Buchak's theory as the von Neumann-Morgenstern theory stands to Savage's; that is, it assumes that the probabilities are given exogenously, as objective chances, perhaps, while Buchak's assumes they are given endogenously, as the individual's credences.

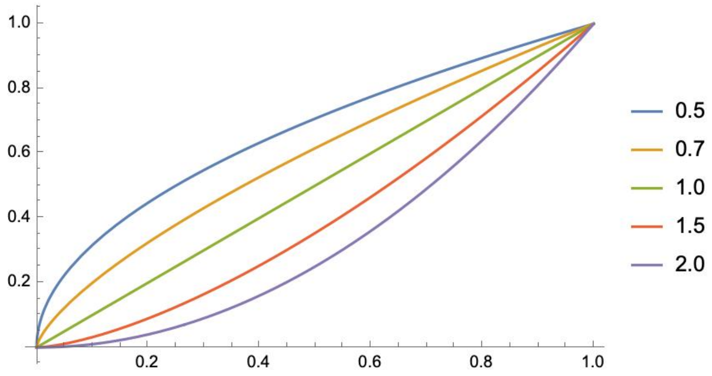

Figure 2: For $m = 0.5, 0.7, 1, 1.5, 2$, we plot $R_m(p) = p^m$ for $0 \leq p \leq 1$.

them as before to give $P^\star_{\mu,s}$. Now we consider two risk functions, R and R'. We ask: what is the risk-weighted expected utility of maximising risk-weighted expected utility with R' and P from the point of view of R and $P^\star_{\mu,s}$? And: for which R is there some alternative R' such that

$$\mathrm{RExp}_{R,P^\star_{\mu,s}}(REU_{R',P}) > \mathrm{RExp}_{R,P^\star_{\mu,s}}(REU_{R,P})$$

Figure 3 gives some answers. In it, we let $p = 0.3$ and write R_m for the risk function $R_m(p) = p^m$. Then, for $m = 0.5, 0.7, 1, 1.5, 2$, we plot

$$\mathrm{RExp}_{R_m,P^\star_{\mu,s}}(REU_{R_k,P}) - \mathrm{RExp}_{R_m,P^\star_{\mu,s}}(REU_{R_m,P})$$

That is, we plot the extent to which R_m judges R_k to perform better than it judges itself to perform. In Figure 2, we plot these risk functions themselves.

Let's work through the lessons of Figure 3:

- R_1 represents risk-neutrality. Risk-weighted expected utility with risk function R_1 is just expected utility theory. Since expected utility theory is self-recommending, as we saw in Section 2, it is no surprise that R_1 takes no other risk function to perform better than it takes itself to perform.

- $R_{0.5}$ and $R_{0.7}$ both represent risk-seeking attitudes. Indeed, for any $0 < m < 1$, R_m is risk-seeking. The lower m, the more weight it requires you to place on the best-case outcome, and the more risk-seeking it is. $R_{0.5}$ takes R_m to perform better than it takes itself to perform for $0.5 < m < 1.1$; $R_{0.7}$ takes R_m to outperform itself for $0.7 < m < 1.03$. So each takes less extremely risk-seeking attitudes to perform better; and each takes R_1, or expected utility theory, to do better.

- $R_{1.5}$ and R_2 both represent risk-averse attitudes. Indeed, for any $1 < m$, R_m is risk-averse. The greater m, the more weight it requires you to place on the worst-case outcome, and the more risk-averse it is. R_2 takes R_m to perform better for $1.2 < k < 2$; $R_{1.5}$ takes R_m to outperform itself for $1.05 < m < 1.5$. So each takes less extremely risk-averse attitudes to perform better; and neither takes R_1, or expected utility theory, to do better.

The upshot: Buchak's theory is not self-recommending in general. There are risk functions R that the theory permits such that REU with R is not self-recommending. But perhaps there

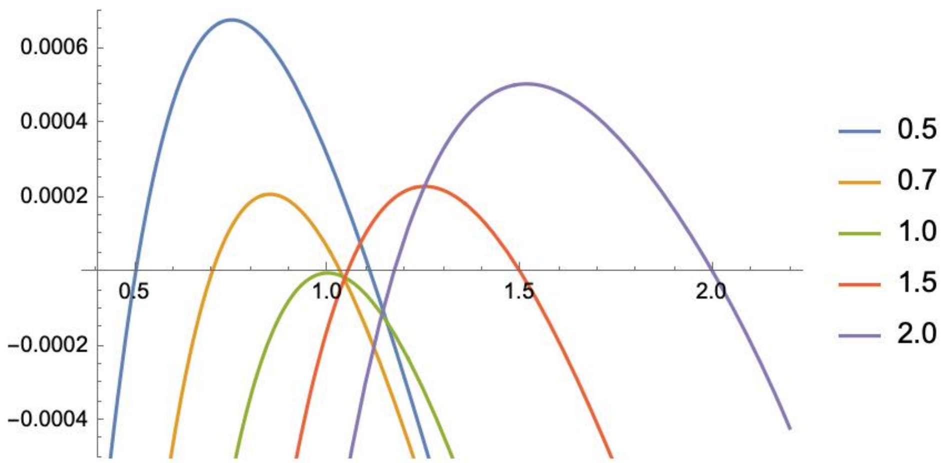

Figure 3: Let $p = 0.3$. Then, for $m = 0.5, 0.7, 1, 1.5, 2$, we plot $\text{RExp}_{R_m, P^\star_{\mu,s}}(REU_{R_k,P}) - \text{RExp}_{R_m, P^\star_{\mu,s}}(REU_{R_m,P})$ over the range $0.4 \leq k \leq 2.1$.

are risk functions other than R_1 that, in combination with REU, give a self-recommending theory? If there were, this would raise an interesting possibility: REU in general is self-defeating; but there could be a more restrictive version, which limits the permissible risk functions an agent may use, and each instance of that version might be self-recommending. That is, it's possible that this argument doesn't rule out REU, but rather limits the permissible risk functions. It's possible. I have no proof that this isn't true. But I do know that, for $m = 0.5, 0.6, \ldots, 0.9$ and for $m = 1.1, 1.2, \ldots, 2$, REU with R_m is not self-recommending.

3.3 Going further

Two assumptions stand out in what I said in the previous section: the assumption that we have a uniform distribution over the possible decision problems we might face; and the assumption that our risk function takes the form $R_m(p) = p^m$ for some m. Because of the computational complexity of calculating $\text{RExp}_{R_m, P^\star_{\mu,s}}(REU_{R_k,P})$, it's difficult to see how things go in general when we remove those assumptions. So, in this section, I'll mention some results that hold when you know a little more about the decision problem you'll face. In particular, I'll assume you know you'll face the decision whether or not to pay a particular price for a bet that pays out a single unit of utility if we're in state s_1 and nothing if we're in state s_2. Let t be the price in question, given in units of utility. Then the decision problem has the form $(0, 0; 1 - t, -t)$. The first option, on which you refuse to pay price t for the bet, gives you 0 in both states; the second option, on which you do pay t, gives you $1 - t$ in state s_1, since you win 1, but you paid t in the first place, and $-t$ in state s_2, since you win nothing and paid t.

First, let's consider a risk function that doesn't have the form $R_m(p) = p^m$. Here's a family of such functions, some of which are illustrated in Figure 4:

$$R^\dagger_n(p) = \frac{p^n}{p^n + (1-p)^n}$$

Then we can see from Figures 5 to 8 that, for many values of n and for μ the uniform dis-

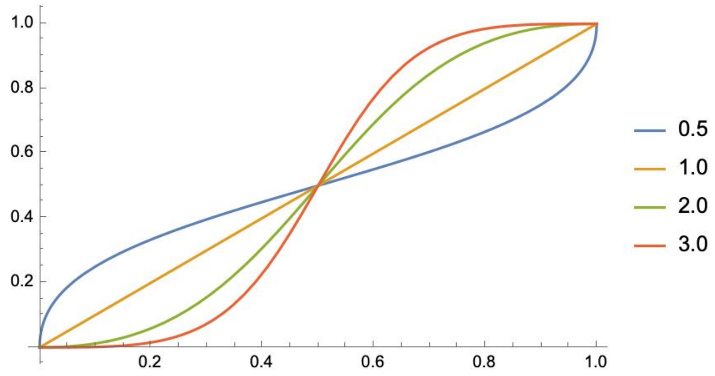

Figure 4: For $n = 0.5, 1, 2, 3$, we plot $R^{\dagger}_n(p)$ for $0 \leq p \leq 1$.

tribution over t, $R^{\dagger}_n$ does not recommend itself. Indeed, we encounter again the same phenomenon we saw in the previous section:

- $R^{\dagger}_1(p) = p$, and so risk-weighted expected utility theory with that risk function is just expected utility theory, and so it recommends itself;
- $R^{\dagger}_{0.5}$ recommends $R^{\dagger}_n$ for $0.5 < n$; that is, it recommends risk functions that are closer to the neutral risk function $R^{\dagger}_1$; and
- $R^{\dagger}_2$ recommends $R^{\dagger}_n$ for $n < 2$; that is, it too recommends risk functions that are closer to $R^{\dagger}_1$; and similarly for $R^{\dagger}_3$.

Next, we return to using risk functions of the form $R_m(p) = p^m$, but this time we use beta distributions over the price t you might pay for the bet. Figure 9 plots some of the probability density functions of the different beta distributions we'll consider. Figure 10 shows that REU with R_2 prefers REU with $R_{2.1}$ over itself for a number of different beta distributions.

4 Decision theories for imprecise credences

4.1 What are imprecise credences?

In both expected utility theory and Buchak's risk-weighted alternative, we represent an individual as assigning precise credences to the various states of the world. But some philosophers, computer scientists, and engineers think that we do better to model individuals as having imprecise credences instead (Walley, 1991; Bradley, 2016). There are many ways to do this, but one of the most well-known ways is to represent an individual's doxastic state not by a single credence function, which assigns to each state of the world a single numerical measure of their confidence in that state, but by a set of such credence functions. So, if the states of the world are just s_1 and s_2, as they have been in the examples we've been considering, we would represent you by a set of probability functions defined on those two states. Following van Fraassen (1990), we call this set your *representor*.

Even within this single proposal, there are multiple ways in which we might interpret the formalism. Here's a natural one: your actual credal state is indeterminate and the credence functions in your representor are its legitimate precisifications. What is determinately true of your credence function is just whatever is true of all the credence functions in your

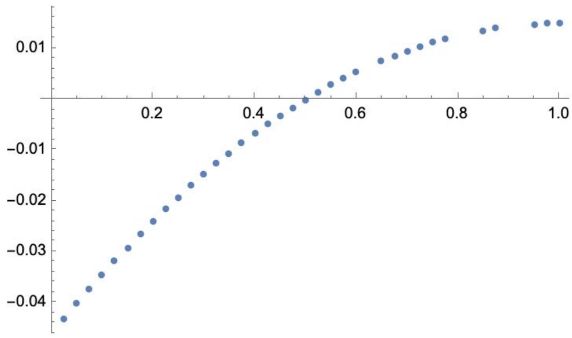

Figure 5: This plots $\mathrm{RExp}_{R^\dagger_{0.5},P^\star_{\mu,s}}(\mathrm{REU}_{R^\dagger_k,P}) - \mathrm{RExp}_{R^\dagger_{0.5},P^\star_{\mu,s}}(\mathrm{REU}_{R^\dagger_{0.5},P})$ for $p = 0.2$, μ the uniform distribution, s chosen arbitrarily, and $0 \leq k \leq 1$.

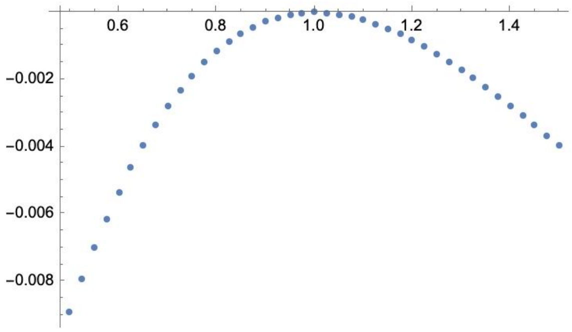

Figure 6: This plots $\mathrm{RExp}_{R^\dagger_{1},P^\star_{\mu,s}}(\mathrm{REU}_{R^\dagger_k,P}) - \mathrm{RExp}_{R^\dagger_{1},P^\star_{\mu,s}}(\mathrm{REU}_{R^\dagger_{1},P})$ for $p = 0.2$, μ the uniform distribution, s chosen arbitrarily, and $0.5 \leq k \leq 1.5$.

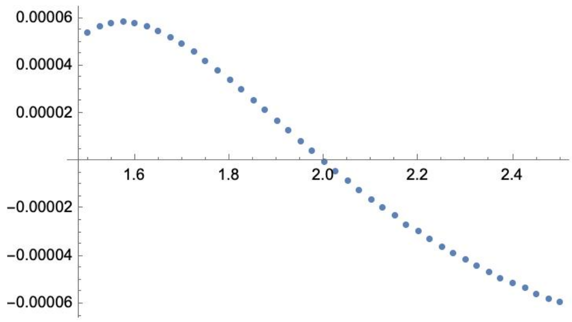

Figure 7: This plots $\mathrm{RExp}_{R^\dagger_{2},P^\star_{\mu,s}}(\mathrm{REU}_{R^\dagger_k,P}) - \mathrm{RExp}_{R^\dagger_{2},P^\star_{\mu,s}}(\mathrm{REU}_{R^\dagger_{2},P})$ for $p = 0.2$, μ the uniform distribution, s chosen arbitrarily, and $1.5 \leq k \leq 2.5$.

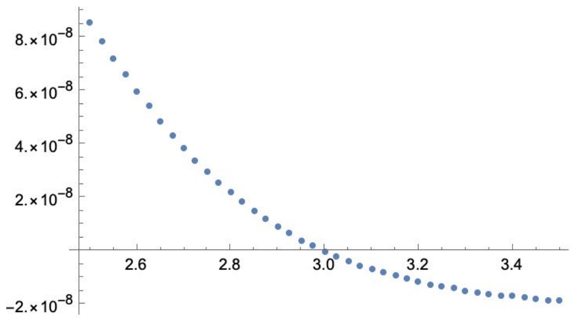

Figure 8: This plots $\mathrm{RExp}_{R^\dagger_{3},P^\star_{\mu,s}}(\mathrm{REU}_{R^\dagger_k,P}) - \mathrm{RExp}_{R^\dagger_{3},P^\star_{\mu,s}}(\mathrm{REU}_{R^\dagger_{3},P})$ for $p = 0.2$, μ the uniform distribution, s chosen arbitrarily, and $2.5 \leq k \leq 3.5$.

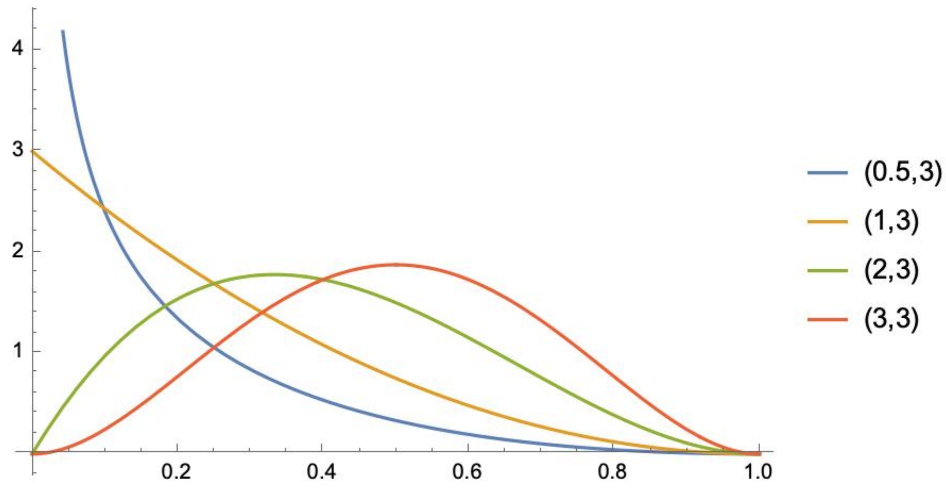

Figure 9: This plots the probability density functions for $B(0.5, 3)$, $B(1, 3)$, $B(2, 3)$, and $B(3, 3)$.

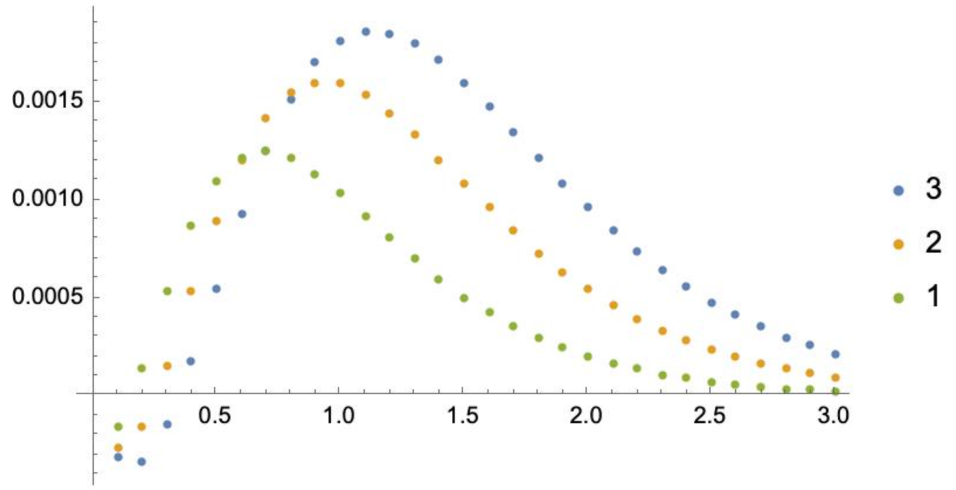

Figure 10: For $\beta = 1, 2, 3$, this plots $\mathrm{RExp}_{R_2, P^\star_{B(\alpha,\beta),s}}(\mathrm{REU}_{R_{2.1},P}) - \mathrm{RExp}_{R_2, P^\star_{B(\alpha,\beta),s}}(\mathrm{REU}_{R_2,P})$ for $p = 0.3$, s chosen arbitrarily, and $0.5 \leq \alpha \leq 3$.

representor. So, for instance, it might be determinately true that your credence that it will rain is greater than 80% but for no value greater than 80% is it determinately true that it is your credence in rain. That would be true if your representor were the set of all probability functions that assign to rain a credence greater than 80%.

In this section, we'll again consider only decision problems in which there are two options, o_1 and o_2, and each option is defined on only two states of the world, s_1 and s_2. And we'll assume that your imprecise credences have a particular form; we'll assume your representor is determined by a pair of probabilities $0 \leq p \leq q \leq 1$ as follows, where we write $(r, 1-r)$ for the probability function that assigns r to s_1 and $1-r$ to s_2:

$$\mathbf{P}_{[p,q]} = \{(r, 1-r) : p \leq r \leq q\}$$

That is, it is determinately true that your credence in s_1 is at least p and at most q, but nothing more specific is determinately true of it.

4.2 What are the decision theories for imprecise credences?

If we represent our decision-maker's doxastic state in this way, what are the possible decision theories they might adopt? I'll try to give a reasonably comprehensive account of what has been proposed.[6] From an informal survey of those who favour this approach, it seems that E-Admissibility and Maximality are the most popular. Unfortunately, they're not self-recommending. Indeed, among the decision theories for imprecise credences I consider here, only one is self-recommending for any representor of the form, $\mathbf{P}_{[p,q]}$; and, what's more, according to that decision theory, an individual with such a representor should behave as if she had precise credences, and in particular, the precise credences that sit at the midpoint of her representor.

Midpoint Expected Utility Theory$_{[p,q]}$ (MEU$_{[p,q]}$ or EU$_{\frac{p+q}{2}}$) Given an option o in

[6]My inventory relies heavily on (Bradley, 2016), (Weatherson, 1998), and an impromptu tutorial by XXX.

$\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that $\text{Exp}_{\frac{p+q}{2}}(o) < \text{Exp}_{\frac{p+q}{2}}(o')$.

This is not popular among those who favour imprecise credences because, for all practical purposes, having imprecise credences and using this rule is indistinguishable from having certain precise ones and using expected utility theory. However, it follows immediately from Theorem 1 that it is self-recommending, since it is simply a version of expected utility theory.

Next, we define a string of theories that appeal to the minimal and maximal expected utility of an option assigned by credences in the representor. To state them succinctly, let's say:

$$l_o = \min_{P \in \mathbf{P}} \text{Exp}_P(o) \qquad \text{and} \qquad m_o = \max_{P \in \mathbf{P}} \text{Exp}_P(o)$$

Global Dominance$_{[p,q]}$ **(GD$_{[p,q]}$)** Given an option o in $\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that $u_o < l_{o'}$.

Γ-Maximin$_{[p,q]}$ **(ΓMaximin$_{[p,q]}$)** Given an option o in $\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that $l_o < l_{o'}$.

Γ-Maxi$_{[p,q]}$ **(ΓMaxi$_{[p,q]}$)** Given an option o in $\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that:

(i) $l_o < l_{o'}$ and $u_o < u_{o'}$,

(ii) $l_o < l_{o'}$ and $u_o = u_{o'}$,

(iii) $l_o = l_{o'}$ and $u_o < u_{o'}$.

Global Dominance is very permissive: an option is only ruled out if there's an alternative such that every expected utility that the alternative receives from a probability function in the representor is higher than every expected utility that original option receives from a probability function in the representor, and that's a lot to ask. Γ-Maximin is a risk-averse decision rule, focussing only on the minimal expected utility that options receive from probability functions in the representor. And Γ-Maxi says that an option is ruled out if there's an alternative with better minimal and maximal expected utilities.

Finally, here are two decision theories that are based on a supervaluationist interpretation of the representor: that is, they rule out options on the grounds that all probability functions in the representor would rule them out.

E-Admissibility$_{[p,q]}$ **(EAd$_{[p,q]}$)** Given an option o in $\mathcal{O}$, o is impermissible iff, for all P in $\mathbf{P}$, there is o' in $\mathcal{O}$ such that,

$$\text{Exp}_P(o) < \text{Exp}_P(o')$$

Maximality$_{[p,q]}$ **(Maximality$_{[p,q]}$)** Given an option o in $\mathcal{O}$, o is impermissible iff there is o' in $\mathcal{O}$ such that, for all P in $\mathbf{P}$,

$$\text{Exp}_P(o) < \text{Exp}_P(o')$$

They differ only in the order of the quantifiers. E-Admissibility rules out an option if it is determinately the case that there is a better alternative, while Maximality rules it out if there is an alternative that is determinately better. Usually, Maximality is more permissive than E-Admissibility, since it requires more to rule something out. But in the cases that concern us, where there are just two options available in our decision problems, they coincide.

4.3 Are these theories self-recommending?

Of the theories described in the previous section—Midpoint Expected Utility, Global Dominance, Γ-Maximin, Γ-Maxi, Maximality, and E-Admissibility—only the first is self-recommending.

As before, we focus on the case in which you know the decision problem you'll will have the form $(a, b; c, d)$, for $0 \leq a, b, c, d \leq 1$. Let's assume your representor is $\mathbf{P}_{[p,q]}$. That is, it contains all and only the probability functions P, with $P(s_1) = r$ and $P(s_2) = 1 - r$, where $p \leq r \leq q$. Now let's extend this representor to a new representor $\mathbf{P}^\star$ that is defined over the meta states of the world. As before, we'll assume that the decision problem you'll face, the state of the world, and the tie-breaker proposition are all independent of one another. And we'll assume that the credences in s_1 and s_2 cover the same range as in $\mathbf{P}$. We'll assume the credences in the tie-breaker propositions lie within a particular range $[t, t']$, though it turns out that it won't matter which range we choose. And we'll assume that our credences over the various decision problems are given by the uniform distribution μ. So

$$\mathbf{P}^\star_{\mu,[t,t']} = \{P^\star_{\mu,s} : t \leq s \leq t'\}$$

where, recall from above, for any measurable subset $\mathcal{D}'$ of $\mathcal{D}$, then

$$P^\star_{\mu,s}(D \in \mathcal{D}' \;\&\; s_1 \;\&\; t_1) = s \times p \times \mu(D \in \mathcal{D}')$$

Let's now work through the decision theories enumerated in the previous section.

4.3.1 Midpoint Expected Utility

It is clear that Midpoint Expected Utility is self-recommending, since it is just a version of expected utility theory, which we saw above is self-recommending.

4.3.2 Global Dominance

To see that Global Dominance is sometimes not self-recommending, consider Figure 11. There, we plot four functions. For each representor $\mathbf{P}_{[p,q]}$ with $0 \leq p \leq 0.4$, $q = 1$, and s taking any value from 0 to 1 inclusive, we plot: (i) $\text{Exp}_{p,s}(\text{MEU}_{[p,q]})$, (ii) $\text{Exp}_{q,s}(\text{MEU}_{[p,q]})$, (iii) $\text{Exp}_{p,s}(\text{GD}_{[p,q]})$, (iv) $\text{Exp}_{q,s}(\text{GD}_{[p,q]})$. Here, for a decision theory DT, $\text{Exp}_{p,s}(\text{DT})$ is the expectation of its utility from the point of view of the probability function $P^\star_{\mu,s}$ in $\mathbf{P}^\star_{\mu,[t,t']}$ such that

$$P(s_1) = p \qquad P(s_2) = 1 - p \qquad P(t_1) = s \qquad P(t_2) = 1 - s$$

Now, you'll notice that we don't actually specify s in this figure. The reason is that the expectation is the same for any s. What the figure shows is that, for $0 \leq p \leq 0.25$ and $q = 1$, the maximum expected utility of Global Dominance (the red line of dots) with representor $\mathbf{P}_{[p,q]}$ lies below the minimum expected utility of Midpoint Expected Utility (the blue line of dots) with the same representor. And so the decision theory of Global Dominance declares itself impermissible. And indeed this is true of representors of other forms as well, such as $P_{[0.1,0.9]}$. However, as we also see, this isn't always the case: for $p = 0$ and q greater than 0.3, for instance, the maximum expected utility of Global Dominance exceeds the minimum expected utility of Midpoint Expected Utility. Of course, that doesn't show that Global Dominance is self-recommending for those representors, since there might be a different decision theory whose minimum exceeds its maximum for those; but it does show that Midpoint Expected

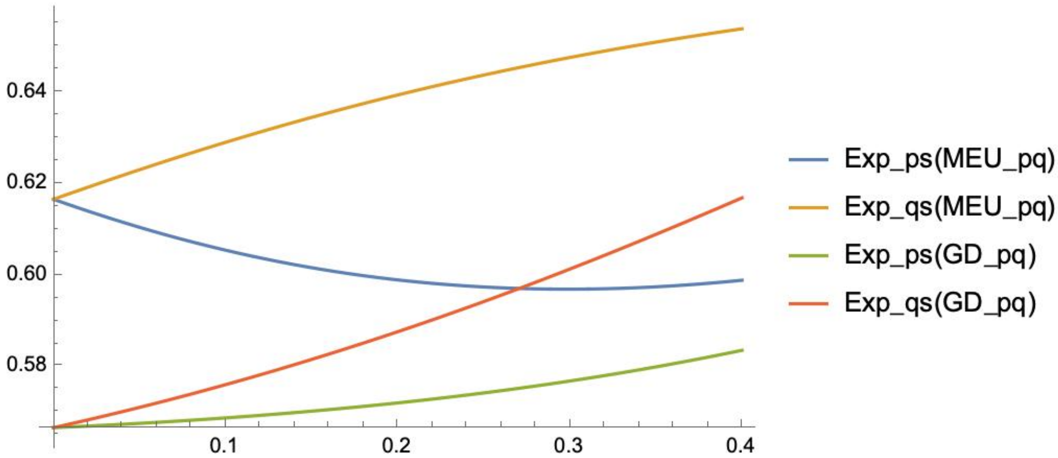

Figure 11: This plots $\mathrm{Exp}_{p,s}(\mathrm{MEU}_{[p,q]})$, $\mathrm{Exp}_{q,s}(\mathrm{MEU}_{[p,q]})$, $\mathrm{Exp}_{p,s}(\mathrm{GD}_{[p,q]})$, $\mathrm{Exp}_{p,s}(\mathrm{GD}_{[p,q]})$, for $0 \leq p \leq 0.4$, $q = 1$, and any value of s between 0 and 1.

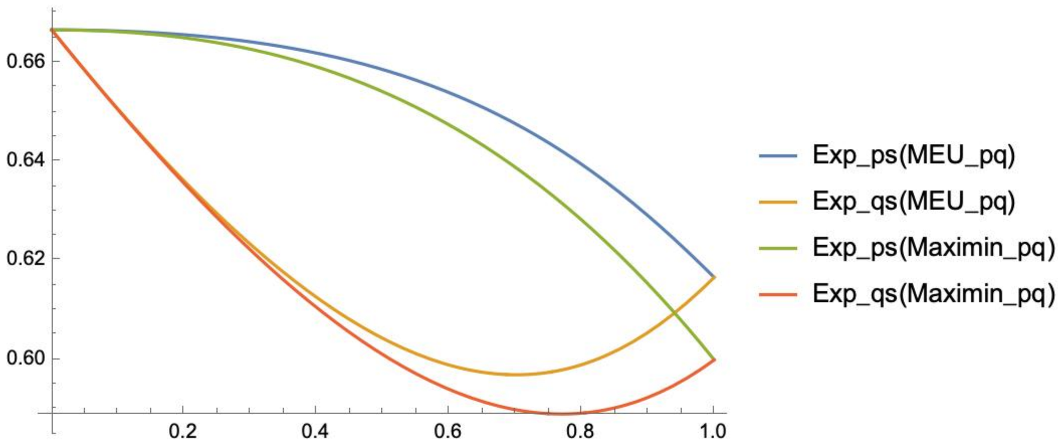

Figure 12: This plots $\mathrm{Exp}_{p,s}(\mathrm{MEU}_{[p,q]})$, $\mathrm{Exp}_{q,s}(\mathrm{MEU}_{[p,q]})$, $\mathrm{Exp}_{p,s}(\Gamma\mathrm{Maximin}_{[p,q]})$, $\mathrm{Exp}_{p,s}(\Gamma\mathrm{Maximin}_{[p,q]})$, for $p = 0$, $0 \leq q \leq 1$, and any value of s between 0 and 1.

Utility doesn't witness Global Dominance's failure to self-recommend in those cases. But the important upshot is that there is some reasonably substantial class of representors for which Global Dominance is not self-recommending.

4.3.3 Γ-Maximin

To see that Γ-Maximin is sometimes not self-recommending, consider Figure 12. There, as for Global Dominance, we plot four functions. For each representor $\mathbf{P}_{[p,q]}$ with $p = 0$ and $0 \leq q \leq 1$, we plot: (i) $\mathrm{Exp}_{p,s}(\mathrm{MEU}_{[p,q]})$, (ii) $\mathrm{Exp}_{q,s}(\mathrm{MEU}_{[p,q]})$, (iii) $\mathrm{Exp}_{p,s}(\Gamma\mathrm{Maximin}_{[p,q]})$, (iv) $\mathrm{Exp}_{q,s}(\Gamma\mathrm{Maximin}_{[p,q]})$. As before, the value of s makes no difference. As we can see, for all the representors of this form, Γ-Maximin is not self-recommending. The minimum expected value of Γ-Maximin is plotted as the red dots, while the minimum expected value of Midpoint Expected Utility is plotted as the yellow dots, and the latter always exceed the former. What's more, there are other representors not of this form for which Γ-Maximin is

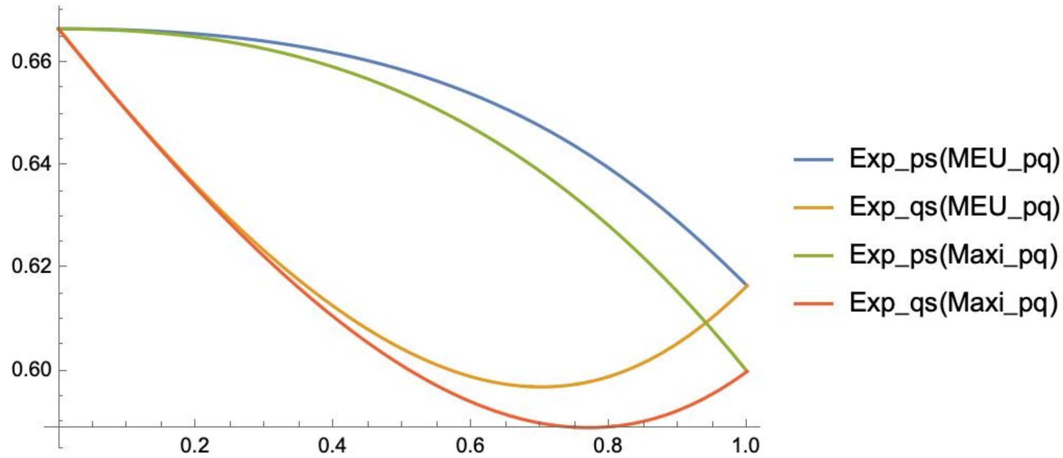

Figure 13: This plots $\text{Exp}_{p,s}(\text{MEU}_{[p,q]})$, $\text{Exp}_{q,s}(\text{MEU}_{[p,q]})$, $\text{Exp}_{p,s}(\Gamma\text{Maxi}_{[p,q]})$, $\text{Exp}_{p,s}(\Gamma\text{Maxi}_{[p,q]})$, for $p = 0$, $0 \leq q \leq 1$, and any value of s between 0 and 1.

not self-recommending: $\mathbf{P}_{[0.1,0.3]}$ is an example.

4.3.4 Γ-Maxi

Our strategy for Γ-Maxi is exactly the same as for Global Dominance and Γ-Maximin. Figure 13 is the relevant figure. For all the representors of this form, Γ-Maxi is not self-recommending. The minimum expected utility of Γ-Maxi is given by the red dots, and the minimum expected utility of Midpoint Expected Utility is given by the yellow ones; and the maximum expected utility of Γ-Maxi is given by the green dots, while the maximum expected utility of Midpoint Expected Utility is given by the blue ones. And again there are other representors for which Γ-Maxi is not self-recommending: again, $\mathbf{P}_{[0.1,0.3]}$ is an instance.

4.3.5 E-Admissibility

The same strategy applies to E-Admissibility as well. Figures 14 to 17 shows how Midpoint Expected Utility and E-Admissibility compare for representors of the form $\mathbf{P}_{[p,q]}$, where $p = 0, 0.25, 0.5, 0.75$, and $p \leq 1$. Note that the range covered by the x-axis is different for each plot.

For all representors pictured, E-Admissibility is not self-recommending. The expected utility of E-Admissibility from the point of view of p is less than the expected utility of Midpoint Expected Utility from that point of view; and similarly for q; and therefore also for any $p \leq r \leq q$. So every probability function in the representor, judges Midpoint Expected Utility better than E-Admissibility. So E-Admissibility is not self-recommending.

4.3.6 Maximality

Finally, Maximality. Here, we don't need a new figure. After all, when faced with binary decision problems, E-Admissibility and Maximality always rule the same options permissible. So Figures 14 to 17 give the expected utilities of using that decision theory. And,

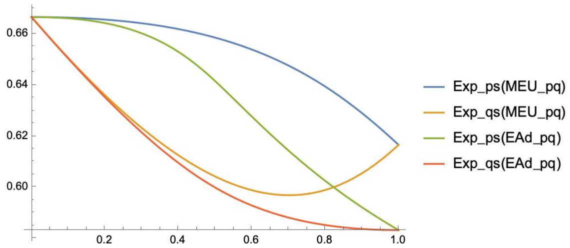

Figure 14: This plots $\text{Exp}_{p,s}(\text{MEU}_{p,q})$, $\text{Exp}_{q,s}(\text{MEU}_{p,q})$, $\text{Exp}_{p,s}(\text{EAd}_{p,q})$, $\text{Exp}_{p,s}(\text{EAd}_{p,q})$, for $p = 0$, $0 \leq q \leq 1$, and any value of s between 0 and 1.

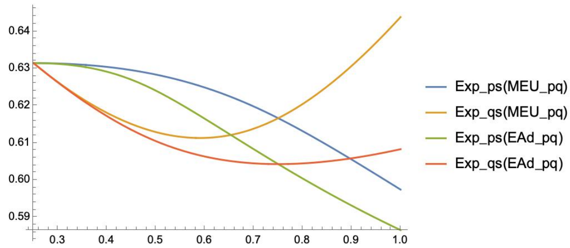

Figure 15: This plots $\text{Exp}_{p,s}(\text{MEU}_{p,q})$, $\text{Exp}_{q,s}(\text{MEU}_{p,q})$, $\text{Exp}_{p,s}(\text{EAd}_{p,q})$, $\text{Exp}_{p,s}(\text{EAd}_{p,q})$, for $p = 0.25, 0.25 \leq q \leq 1$, and any value of s between 0 and 1.

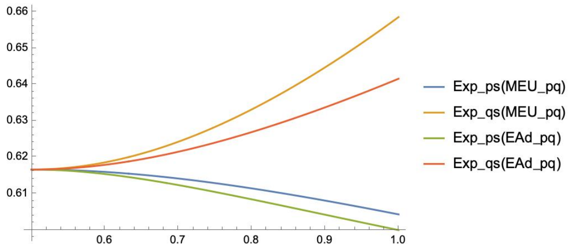

Figure 16: This plots $\text{Exp}_{p,s}(\text{MEU}_{p,q})$, $\text{Exp}_{q,s}(\text{MEU}_{p,q})$, $\text{Exp}_{p,s}(\text{EAd}_{p,q})$, $\text{Exp}_{p,s}(\text{EAd}_{p,q})$, for $p = 0.5, 0.5 \leq q \leq 1$, and any value of s between 0 and 1.

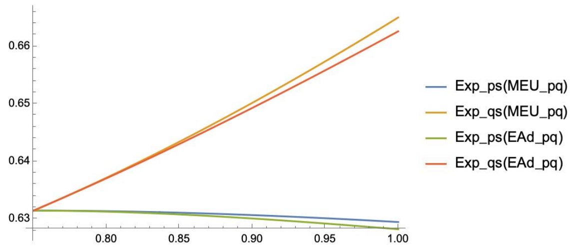

Figure 17: This plots $\text{Exp}_{p,s}(\text{MEU}_{p,q})$, $\text{Exp}_{q,s}(\text{MEU}_{p,q})$, $\text{Exp}_{p,s}(\text{EAd}_{p,q})$, $\text{Exp}_{p,s}(\text{EAd}_{p,q})$, for $p = 0.75, 0.75 \leq q \leq 1$, and any value of s between 0 and 1.

what's more, it shows that Maximality is not self-recommending, for there is a decision theory, namely, Midpoint Expected Utility, that every probability function in the representor expects to outperform Maximality.

4.4 Going further

Two notable assumptions from this section are that the decision problems we might face will have just two options, and that those options will be defined on just two states of the world. While the computational complexity prevents us from giving fully general answers when we lift these restrictions, we can offer some results.

First, suppose we know we will face a decision problem of the following form:

	s_1	s_2	s_3
o_1	a	b	c
o_2	d	e	f

which we might write as $(a, b, c; d, e, f)$. Then, consider the following representor:

$$\mathbf{P} = \{\lambda P + (1 - \lambda)P' : 0 \leq \lambda \leq 1\}$$

where P and P' are probability functions defined on the three states s_1, s_2, s_3 as follows:

	s_1	s_2	s_3
P	0.1	0.2	0.7
P'	0.3	0.2	0.5

Then each of the decision rules other than Midpoint Expected Utility prefers to use Midpoint Expected Utility with $\mathbf{P}$ to using itself with $\mathbf{P}$. That is, the minimal expected utility of Γ-Maximin with $\mathbf{P}$ is less than the minimal expected utility of Midpoint Expected Utility with $\mathbf{P}$, each probability function in $\mathbf{P}$ expects Midpoint Expected Utility with $\mathbf{P}$ to perform better as a decision theory than it expects E-Admissibility or Maximality to perform, and so on.

Second, suppose we know we will face a decision problem of the following form:

	s_1	s_2
o_1	a	b
o_2	c	d
o_3	e	f

where we might write $(a, b; c, d; e, f)$. Then consider the following representor:

$$\mathbf{P} = \{\lambda P + (1 - \lambda)P' : 0 \leq \lambda \leq 1\}$$

where P and P' are probability functions defined on the three states s_1 and s_2 as follows:

	s_1	s_2
P	0.3	0.7
P'	0.6	0.4

Then again each of the decision rules other than Midpoint Expected Utility prefers to use Midpoint Expected Utility with $\mathbf{P}$ to using itself with $\mathbf{P}$. An interesting side note: when there are three options, E-Admissibility and Maximality are no longer equivalent, for there can be decision problems in which each probability function in the representor expects the second or third option to be better than the first, but it's not the case that each probability function expects the second to be better than the first and it's not the case that each expects the third to be better than the first. And it turns out that, given the representor we've just described, Maximality actually prefers E-Admissibility to itself, and each prefers Midpoint Expected Utility to itself: that is, every probability function in the representor expects you do to better by using E-Admissibility than by using Maximality.

The results of this section are troubling for those who would represent our doxastic state using a set of probability functions instead of a single one. It is often pointed out that Γ-Maximin is diachronically incoherent. But, as I argued above, such an objection has little bite. What tells much more strongly against the view is that each of the existing decision theories for imprecise probabilities judge Midpoint Expected Utility to be better as a means to your ends than they judge themselves to be. That is, by their own lights, they are impermissible.

5 Decision theories for total ignorance

The final set of decision theories I will consider were developed initially at the Cowles Commission in Chicago in the 1950s by a group of young researchers who would go on to make significant contributions in other areas of their subjects: the mathematicians Abraham Wald and John Milnor, and the economists Leonard Hurwicz and Kenneth Arrow. They were interested in how we should make decisions in those situations in which our uncertainty is so profound that we are unable to assign probabilities, even subjective ones; that is, situations of so-called Knightian uncertainty. Here are three of the decision rules they formulated:

Maximin Given an option o in $\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that

$$\min_{s \in \mathcal{S}} o(s) < \min_{s \in \mathcal{S}} o'(s)$$

Maximax Given an option o in $\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that

$$\max_{s \in \mathcal{S}} o(s) < \max_{s \in \mathcal{S}} o'(s)$$

Hurwicz Criterion$_\lambda$ Given an option o in $\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that

$$\lambda \max_{s \in \mathcal{S}} o(s) + (1 - \lambda) \min_{s \in \mathcal{S}} o(s) < \lambda \max_{s \in \mathcal{S}} o'(s) + (1 - \lambda) \min_{s \in \mathcal{S}} o'(s)$$

where $0 \leq \lambda \leq 1$.

Note that:

- Hurwicz$_0$ Criterion = Maximin
- Hurwicz$_1$ Criterion = Maximax

Now we can ask whether these decision theories are self-recommending or not. It will turn out that all of them are.

I will consider Maximin first, and pretty much the same reasoning can be applied to Maximax and any version of the Hurwicz Criterion. As above, we'll consider any two-state, two-option decision problem in which the utilities of the options are bounded above and below. So we're considering any decision problem of the form $(a, b; c, d)$, where $0 < a, b, c, d < 1$. Now suppose that we know we'll face such a decision problem in the future, but we don't know which. Then let us ask, from the point of view of Maximin itself, how Maximin will perform if used to make this decision. To do this, we begin by asking what is the lowest utility it might obtain for us? And we note that it is 0, since we might face a decision of the form $(0, b; 0, d)$. But then we note that the same is true of any decision theory: if you might face $(0, b; 0, d)$ at state s_1, then your decision theory will obtain for you 0 for sure. And so Maximin is permissible by its own lights—that is, it is self-recommending.

However, we're not done, for we know that Maximin has a different flaw. It permits options that are weakly dominated by others. For instance, faced with $(0, \frac{1}{2}; 0, 1)$ it permits us to choose the first option when the second weakly dominates it. Indeed, this fact makes it rather less impressive that it is self-recommending. After all, one way in which a decision theory can ensure that it is self-recommending is to be very permissive—if it doesn't rule much out, it won't rule itself out. This is why it is surprising that expected utility theory is self-recommending, but Global Dominance isn't. Expected utility theory is quite strict, while Global Dominance is quite lenient. Nonetheless, we can easily tweak Maximin to avoid this flaw:

Maximin$^+$ Given an option o in $\mathcal{O}$, o is impermissible iff there is an alternative option o' in $\mathcal{O}$ such that

(i) $\min_{s \in \mathcal{S}} o(s) < \min_{s \in \mathcal{S}} o'(s)$, or

(ii) $o(s) \leq o'(s)$ for all s in $\mathcal{S}$, and $o(s) < o'(s)$ for some s in $\mathcal{S}$.

And now we can ask whether this less permissive variant, Maximin$^+$, is self-recommending. Again, the answer is yes. As before, Maximin$^+$, like all decision theories, gives utility 0 in the worst-case scenario. But what's more, Maximin$^+$ isn't weakly dominated. Suppose, for a contradiction, that it is. Then there is a decision theory DT such that

$$\text{Maximin}^+(\langle(a,b;c,d),s_i,t_j\rangle) \leq \text{DT}(\langle(a,b;c,d),s_i,t_j\rangle)$$

for all meta states, with strict inequality for some. Then take one of the meta states for which strict inequality holds, so that

$$\text{Maximin}^+(\langle(a^*,b^*;c^*,d^*),s_{i^*},t_{j^*}\rangle) < \text{DT}(\langle(a^*,b^*;c^*,d^*),s_{i^*},t_{j^*}\rangle)$$

Then, for $i = 1, 2$,

$$\text{Maximin}^+(\langle(a^*,b^*;c^*,d^*),s_i,t_{j^*}\rangle) < \text{DT}(\langle(a^*,b^*;c^*,d^*),s_i,t_{j^*}\rangle)$$

with inequality for $i = i^*$. But then whichever option Maximin$^+$ chooses when faced with decision problem $(a^*,b^*;c^*,d^*)$ and tie-breaker t_{j^*}, it is weakly dominated by whichever option DT chooses. But that's a contradiction.

So Maximin$^+$ does not permit weakly dominated options, and it is self-recommending. By similar reasoning, we can tweak Maximax in the same way to avoid weakly dominated options, and the tweaked version will also be self-recommending. And the same goes for the Hurwicz$_\lambda$ criterion.

6 Antecedents to this approach

In Section 1, I noted that exploitability and money pump arguments are among the philosophical antecedents to my approach here; but there are also formal antecedents, and I will describe them here.

The approach is inspired ultimately by I. J. Good's approach to the question of whether we should accept free evidence (Good, 1967). Suppose you are offered some free evidence. Suppose further that you know that, if you accept that evidence, you'll respond to it as the Bayesian demands by conditioning on it. And suppose yet further that you know you'll face a decision after receiving that evidence, you know exactly what the decision is, and you know you'll choose from the options available by maximising expected utility from the perspective of whichever credences you have at that time. Then, Good shows, from the perspective of your current credences, you maximise expected utility by taking the evidence and choosing in line with your new updated credences; refusing the evidence and choosing in line with your current credences has at most the same expected utility, and for some decision problems, the inequality is strict.

Peter M. Brown (1976) then repurposes this result to show that, if you know what decision problem you'll face, and you know you'll receive new evidence, you expect yourself to do at least as well by conditioning on that evidence than by responding in any other way, and for some decision problems you might face, you expect yourself to do better.

Mark Schervish (1989) significantly generalises this approach, and Ben Levinstein (2017) provides a philosophical gloss. Most importantly, they remove the assumption that you know which decision problem you'll face. Instead, they assume that you place a probability distribution over various possibilities. They assume that you know you'll face a decision of the form $(0, t; 1-t, 0)$, where $0 \leq t \leq 1$, and you place a probability distribution over the value of t. As with Good's and Brown's arguments, however, they retain the assumption that you'll choose using expected utility theory. They use this to measure how much utility a particular credence can be expected to obtain for you given a particular way the world is. In this way, Schervish and Levinstein provide a way to score credences, and they show that the scores they generate in this way are the strictly proper scores that are used in judgment elicitation (Savage, 1971) and epistemic utility theory (Predd et al., 2009; Pettigrew, 2016).

Later, Catrin Campbell-Moore and Bernhard Salow (ta) applied this technique under the assumption that we choose using risk-weighted expected utility theory and Jason Konek (2019) applied it in the framework of imprecise credences. In that branch of the literature, we fix the decision theory and use it to provide a way of scoring how well our credences perform when they are fed into it. In the approach I have been pursuing here, we fix the credences and use the approach to assess the decision theories into which we feed them.

7 Conclusion

If we attempt to adjudicate between decision theories by appealing to our intuitions about what rationality requires, we are often led to a stalemate. One way to avoid that is to ask how well the decision theories perform as means to our ends. But we need a theory of means-ends rationality in order to judge that. So we ask each decision theory: would you recommend yourself to me as a good means to acquiring my ends? In this paper, I have shown that many rivals to expected utility theory answer this question in the negative: they are not self-recommending; they are self-undermining. I submit that this rules out those theories as accounts of means-ends rationality.

References

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine. *Econometrica*, *21*(4), 503–546.

Bradley, S. (2016). Imprecise Probabilities. In E. N. Zalta (Ed.) *Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.

Brown, P. M. (1976). Conditionalization and expected utility. *Philosophy of Science*, *43*(3), 415–419.

Buchak, L. (2013). *Risk and Rationality*. Oxford, UK: Oxford University Press.

Buchak, L. (2022). Normative Theories of Rational Choice: Rivals to Expected Utility. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.

Campbell-Moore, C., & Salow, B. (ta). Accurate Updating for the Risk-Sensitive. *British Journal for the Philosophy of Science*.

Christensen, D. (1991). Clever Bookies and Coherent Beliefs. *Philosophical Review*, *100*(2), 229–247.

Gibbard, A., & Harper, W. (1978). Counterfactuals and Two Kinds of Expected Utility. In E. F. M. Clifford Alan Hooker, James L. Leach (Ed.) *Foundations and Applications of Decision Theory*, vol. 13a of *University of Western Ontario Series in Philosophy of Science*, (pp. 125–162). D. Reidel.

Good, I. J. (1967). On the Principle of Total Evidence. *The British Journal for the Philosophy of Science*, *17*, 319–322.

Gustafsson, J. E. (2022). *Money-Pump Arguments*. Cambridge Elements in Decision Theory and Philosophy. Cambridge, UK: Cambridge University Press.

Jeffrey, R. (1965). *The Logic of Decision*. New York: McGraw-Hill.

Joyce, J. M. (1999). *The Foundations of Causal Decision Theory*. Cambridge Studies in Probability, Induction, and Decision Theory. Cambridge: Cambridge University Press.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263.

Konek, J. (2019). IP Scoring Rules: Foundations and Applications. In *Proceedings of the Eleventh International Symposium on Imprecise Probabilities: Theories and Applications*, vol. 103 of *Proceedings of Machine Learning Research*, (pp. 256–264).

Levinstein, B. A. (2017). A Pragmatist's Guide to Epistemic Utility. *Philosophy of Science*, *84*(4), 613–638.

Machina, M. J. (1982). 'Expected Utility' Analysis without the Independence Axiom. *Econometrica*, *50*(2), 277–323.

Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Predd, J., Seiringer, R., Lieb, E. H., Osherson, D., Poor, V., & Kulkarni, S. (2009). Probabilistic Coherence and Proper Scoring Rules. *IEEE Transactions of Information Theory*, *55*(10), 4786–4792.

Quiggin, J. (1982). A Theory of Anticipated Utility. *Journal of Economic Behavior and Organization*, *3*, 323–43.

Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons.

Savage, L. J. (1971). Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, *66*(336), 783–801.

Schervish, M. J. (1989). A general method for comparing probability assessors. *The Annals of Statistics*, *17*, 1856–1879.

Skyrms, B. (1982). Causal Decision Theory. *The Journal of Philosophy*, *79*(11), 695–711.

Stalnaker, R. C. (1972 [1981]). Letter to David Lewis. In W. L. Harper, R. C. Stalnaker, & G. Pearce (Eds.) *Ifs: Conditionals, Belief, Decision, Chance and Time*, University of Western Ontario Series in Philosophy of Science, (pp. 151–2). Dordrecht: Reidel.

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, *23*(3), 165–180.

van Fraassen, B. C. (1990). Figures in a probability landscape. In J. M. Dunn, & A. Gupta (Eds.) *Truth or Consequences*, (pp. 345–56). Dordrecht: Kluwer.

von Neumann, J., & Morgenstern, O. (1947). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press, 2^{nd} ed.

Wald, A. (1945). Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, *46*(2), 265–280.

Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*, vol. 42 of *Monographs on Statistics and Applied Probability*. London: Chapman and Hall.

Weatherson, B. (1998). Decision making with imprecise probabilities. Unpublished manuscript.