Can Imprecise Probabilities be Practically Motivated?

(Draft)

Abstract: It is often claimed that in response to certain types of evidence we are rationally required to adopt belief states that are best represented by imprecise probabilities: sets of probability functions, rather than single ones. Imprecise probabilities, or ranges of probabilities (like 60-80%), have also been advocated in the reporting of scientific information to policy-makers, in engineering, and in the design of artificial intelligence systems. The question motivating this paper is this: should we expect the usage of imprecise probabilities to result in better decision-making than the usage of precise probabilities in these contexts? I will argue that we should not expect imprecise probabilities to do better by undermining the only decision-theoretic motivation for imprecision that has been offered in the literature: ambiguity aversion. If we can't expect imprecision to help us make better decisions, we need to rethink whether there are good reasons to use imprecise probabilities in contexts in which good decision-making is what's of primary concern.

1. Introduction

How confident are you that your neighbor will wear a red shirt tomorrow? You might think that your degree of confidence in this proposition (your credence), won't be represented by a numerical probability like 0.2, 0.21 or 0.210005. Instead, a common way of representing our belief states in such cases makes use of a set of probability functions (or, in the case of a particular proposition, a range of probabilities like [0.1, 0.3]) rather than a single one. These sets of probability functions are called "imprecise" or "mushy" probabilities.

Suppose that Alice and Bob, who share all the relevant evidence, are contemplating whether neighbor Nadeem will wear a red shirt tomorrow. Alice adopts [0.1, 0.3]. But Bob sticks his neck out and adopts the precise probability of 0.2. Here's a question: should Alice expect herself to do better than Bob in virtue of having adopted [0.1, 0.3] rather 0.2?

Let's make the question more precise by getting clear on what's meant by "do better." There are (at least) two important ways in which an agent might be better off as a result of adopting one belief state rather than another: First, the agent might be more *accurate* (an agent is accurate insofar as she has high confidence in truths and low confidence in falsehoods). Second, the agent might do better *practically* – she might be richer, happier, or better achieve her practical aims, whatever they are. For each of these notions of "better off" we can think (or Alice can think) about which of Alice or Bob we expect to do better.

A number of results[1] suggest that imprecise probabilities can't be motivated by accuracy considerations. The purpose of this paper is to see whether imprecise probabilities can be given a practical motivation: whether we can reasonably expect Alice to do better practically than Bob.

Assuming that the world doesn't give out prizes directly for being imprecise, if Alice can be expected to do better than Bob practically, this is because Alice can be expected to make better decisions than Bob. And, indeed, the idea that imprecise probabilities should be used in contexts in which decision-making is of primary concern has been endorsed in a wide range of disciplines spanning from engineering to artificial intelligence to climate science. Here are some representative quotes:

In engineering:

> In situations where structural safety is a major issue, credibility of statistical statements is most crucial because wrong conclusions can be very expensive or even life-threatening. As a consequence, appropriate modeling of uncertainty has become a fundamental challenge…In recent years, imprecise probability models have proven to be a powerful tool for that purpose" (Augustin and Hable (2010), 358).

Aughenbaugh and Paredis (2005)'s paper is aimed at supporting the hypothesis that:

> "in engineering design decisions, it is valuable to explicitly represent the imprecision in the available characterization of uncertainties by using imprecise probabilities" (1).

In artificial intelligence:

> [I]mprecise probability models are needed in inference problems characterized by scarce, vague or conflicting information. These are all characteristics of many real problems, and hence dealing credibly with them seems to be an essential step for realistic decision making. Artificial intelligence is also naturally concerned with decision making, very often in real-world domains. So it is not surprising that artificial intelligence has a long history of research and interest in imprecise probability models (Zaffalon and de Cooman (2005), 1-2).

In climate science:

[1] Seidenfield et al. (2012), Mayo-Wilson and Wheeler (2016) and Schoenfield (2017).

> For climate applications this raises an obvious question: where is the [probability function] going to come from? Any such summary distribution will require us to aggregate information of varying quality from a variety of largely unvalidated models. This will be a heavily subjective exercise – we have low confidence in our ability to discern between models. An alternative to this approach is to recognize that we do not have the quality of information needed to define a unique [probability function] (Heal and Millner (2015), 8).

They go on to propose that "instead of forcing our information into a probabilistic strait-jacket" we should use imprecise probabilities in reporting scientific information to policy-makers.

And finally, Bradley, Hegelson and Hill (2017), in explaining why imprecise probabilities are favored by the Intergovernmental Panel on Climate Change say:

> [one] may be unable to supply the required subjective probabilities and that any "filling in" of the gap between probability ranges and precise probabilities may prove too arbitrary to be a reasonable guide to decision. Policy makers may quite reasonably refuse to base a policy decision on a flimsy information base, especially when there is a lot at stake (505).

My aim is to argue against the view that, given what's sometimes called "ambiguous evidence," using imprecise probabilities will help us make better decisions. In particular, there is nothing that tells in favor of decision theories that make use of imprecise probabilities over a procedure which aggregates the probability functions in the relevant set, and then uses a precise decision theory to make decisions on its basis. If I'm right, then we need to revisit our insistence on the usage of imprecise probabilities in contexts in which decision-making is of primary concern.

My arguments against the benefits of using imprecise probabilities in decision-making will proceed by undermining the only decision theoretic motivation that philosophers, engineers, and climate-scientists have offered: the purported desirability of ambiguity aversion. The seemingly sensible idea behind ambiguity aversion is this: when we don't know the probabilities we should proceed with caution. The thought is that Bob, who acts on his 0.2 probability, will be less cautious than Ann, who uses the entire interval [0.1, 0.3] when making decisions, and that the extra caution that Ann will display in making use of imprecise decision rules is desirable. I will argue otherwise.

Before proceeding, some clarifications about this project:

Clarification 1: If you're familiar with a certain strand in the literature on imprecise probabilities, which focuses on the potentially *disastrous* consequences of moseying around in an imprecise fashion (Seidenfeld (1998) and (2004), Elga (2010), Steele (2010)), the question of whether we can expect there to be any *benefits* from being imprecise might strike you as odd. I'm going to set worries about imprecise probabilities stemming from these potential disasters aside. Maybe there are disastrous consequences for imprecise agents under certain circumstances. But if being imprecise also carries benefits, and the world isn't riddled with evil bookies who know exactly what our belief states are like, or extremely generous ones who are determined to give away their money in ways that only precise agents can enjoy, there may be a trade-off to be had. So I want to approach the question from the glass half-full perspective. In addition to any potential dangers that imprecise agents will face, will being imprecise carry any benefits?

Clarification 2: For the purposes of this paper, I'm going to remain neutral with respect to the question of whether imprecise probabilities are better supported by ambiguous evidence than precise probabilities. If it's true that Alice's attitude is better supported by their shared evidence than Bob's, but we can't expect Alice to make better decisions than Bob, or to be more accurate than Bob, then it will turn out that, insofar as it's desirable to have attitudes well supported by one's evidence, this desirability can't be explained in terms of our interest in accuracy, or in terms of our practical aims. So it's still worth asking, when our sole concern is designing a good spaceship, which sorts of probabilities do we want to input into our decision-making apparatus? This is the sort of question I'll be focusing on. So I won't be arguing that there are *no* motivations for imprecise probabilities (there may be evidential motivations, or perhaps other theoretical motivations). My aim here is just to argue against the view that imprecise probabilities are valuable in decision-making.

Clarification 3: In addition to evaluating a decision procedure in terms of the outcomes it produces, we humans care about such matters as being able to justify the procedure to others, avoiding blame, and so forth. If we think we'll be derided for deciding on the basis of 0.2 but we won't be derided for deciding on the basis of [0.1, 0.3,] then that is of course a consideration that can play a role in choosing a decision procedure. That is not the kind of practical motivation that I'm looking for here, but I'll return to this possibility briefly at the end.

2. Ambiguity Aversion

Suppose I have the option of going to a casino where I bet on propositions about geology or going to a casino where I bet on propositions about the history of philosophy. If my goal is to make money, I'm better off going to the philosophy casino. Why? Because I know more about the history of philosophy than I do about geology. In general, we're better off betting under circumstances in which we have more information than under circumstances in which there is a great deal of uncertainty.

Some people think that what goes for factual uncertainty also goes for what we might call "probabilistic uncertainty." All else equal, the thought goes, we're better off making bets on propositions in which we "know" the probabilities than on propositions whose probabilities are unknown (more on what it might mean to "know" probabilities later). The preference for acting on known rather than unknown probabilities is called "ambiguity aversion."

I'll give a precise definition of ambiguity aversion in a moment, but first I want to present some examples. The first is based on a case from Ellsberg (1961), the fellow who made ambiguity aversion famous:

> URNS: If you draw a red marble from an urn you win a prize. You have a choice between which of two urns to draw from: You can draw from KNOWN, which you know contains 100 marbles: 50 red ones and 50 black ones. Or you can draw from UNKNOWN, an urn that also contains 100 marbles, all of which are red or black, but you don't know what the ratio of red to black marbles is. You choose KNOWN. After all, you think, you at least have a 50% chance of winning if you choose KNOWN, whereas, for all you know, you have no chance of winning if you choose UNKNOWN.

Here's a somewhat less artificial example:

> THE MAYOR: The Town Mayor wants to implement an exciting new program in the city hospital. She meets with her advisory committee to determine whether implementing the program this year is feasible and she is told the following: if it is a dry winter (rainfall $<$ threshold t), then the new program can be implemented unproblematically. However, if funds are used to implement the program and the winter is wet (rainfall $\geq t$), there will be a problem: some of the roads will be in disrepair and the funds needed to fix them won't be available. The mayor asks her meteorologists how likely it is that the winter will be wet and they tell her that the probability of a wet winter is 0.5. The mayor decides to go ahead with the program. If she were told that the probability of a wet winter were anything above 0.5, she wouldn't be willing to take the risk. The next day the meteorologists return: "We're terribly sorry, but we actually made an error in our calculations. All we can tell you is that the probability of a wet winter is between 20% and 80%." The mayor thinks: "In that case, we better not implement the program. After all, the chance of a wet winter could be as high as 80%. It would be irresponsible to take such a risk."

In both cases, the agents have an all-else-equal preference for taking options in which the probabilities of the various outcomes given that option are known. A preference for taking options in which the probabilities of the various outcomes are known turns out to be equivalent to the following, which is how I will be thinking about ambiguity aversion for the purposes of the paper:

> *You're ambiguity averse relative to a set of probability functions P* if, when evaluating an option, you weight the probability functions in *P* on the basis of their degree of pessimism about that option: the more pessimistic a probability function is about an option, the more weight that probability function gets in the evaluation of the option.

Just for the purposes of illustration, let's suppose that the individual probability functions in a given set evaluate an option's desirability by calculating its (classical) expected utility. That is, where *A* is an option, *p* is a probability function, *u* is a utility function, and the O_i are the possible outcomes of taking option *A*, *p* evaluates *A* as follows:

$$EU_p(A) = p(O_1|A)u(O_1) + \ldots p(O_n|A)u(O_n)$$

The most extreme ambiguity averse theory – Γ-maximin[2] – tells us to evaluate each option in accordance with the probability function in the set that is *most* pessimistic about that option: the probability function that assigns that option lowest expected utility. It then recommends that we choose the option that gets the best of these evaluations, that is, the option that has the highest maximally-pessimistic expected utility.

Here's how this works in URNS: suppose we think of the relevant set of probability functions as the set of functions that represent the possible ratios of marbles in the urns. All of these functions assign probability 0.5 to drawing red from KNOWN. This means that if the utility of winning the prize is 1, all of the probability functions will assign an expected utility of 0.5 to drawing from KNOWN. The probability functions in the set will vary anywhere between 0 and 1 with respect to the expected utility of drawing from UNKNOWN. The most pessimistic function will assign expected utility 0 to an UNKNOWN draw. Since 0.5 is greater than 0, Γ-maximin tells us to choose KNOWN over UNKNOWN.

The Γ-maximin treatment of cases like the urn case yields extremely unintuitive results. It tells you, for example, that if you have a ticket that will win you a million dollars if a red marble is drawn from UNKNOWN, you should be willing to trade your ticket for one that will win you one cent if you draw red from KNOWN. But there are less extreme theories which deliver the ambiguity averse verdict in URNS that allow the pessimists to get extra weight, but not all the weight.[3]

Ambiguity aversion is considered by many to be a rational feature of decision-makers. Al-Najjar and Weinstein (2009) say that "a leading interpretation of the ambiguity aversion literature

[2] Such a theory was first proposed and defended by Gärdenfors and Sahlin (1982).
[3] A proposal along these lines is developed in Klibanoff et al. (2005). See also Binmore (2008), Ghiradato et al. (2004) and Gilboa and Marinacci (2011).

is that the [ambiguity averse] choices are *rational* responses by decision-makers to a lack of reliable information that prevents them from forming beliefs with confidence" (252, emphasis in original). Bradley, Hegelson and Hill (2017) suggest that decision theories that do not permit ambiguity aversion have "unintuitive, and indeed normatively undesirable consequences" (504). They claim that "some degree of discrimination between high and low confidence situations does seem appropriate for important policy decisions" (508). Indeed, in the engineering and climate science applications I described above, imprecise probabilities are advocated precisely because it is thought to be desirable to be ambiguity-averse. The desirability of ambiguity aversion has also been appealed to in defending imprecise probabilities against objections like those leveled in Elga (2010) (see e.g. Sahlin and Weirich (2014)).

Ambiguity aversion is generally motivated by appeal to intuitions about cases or by the more general thought that it's appropriate to be cautious when probabilities are unknown. A more systematic defense of ambiguity aversion is offered by Aughenbaugh and Paredis (2005). They perform a computational experiment in which pressure valves are designed in ambiguity averse and non-ambiguity averse ways. The result of the experiment is that, when there is a great deal of probabilistic uncertainty, the ambiguity averse designers do better, on average, than their non-ambiguity averse counterparts.

Below I'll argue that there is nothing desirable about being ambiguity averse, and diagnose the problem with Aughenbaugh and Paredis' study.

3. Against Ambiguity Aversion

There are different strategies one might deploy in arguing against the usage of a decision rule. A common strategy is to show that somebody who uses the rule will, in some sense or another, end up in a bad way. For example, the agent may be susceptible to accepting bets that guarantee a sure loss (a "Dutch Book") or be willing to pay to avoid information, or expect to do worse in the long run than they would if they adopted some other strategy. In fact, any agent that can't be represented as a (classical) expected utility maximizer is susceptible to such consequences.

In particular, risk averse agents – those that, all else equal, prefer options that minimize the spread of utility across possible outcomes – are Dutch-bookable, will pay to avoid information, and will expect to do worse in the long run than a classical expected-utility maximizer.[4] And indeed, some of the literature on ambiguity aversion seems to presuppose that the *only* objection one might have towards ambiguity aversion is that an ambiguity averse agent can't be represented as an expected-utility maximizer. But that is not my objection. It is important to me, therefore, that the arguments I give are not arguments against *any* decision theory that is incompatible with being

[4] See e.g. Buchak (2010) and Buchak (2013) Ch.7.

representable as an expected utility maximizer. For while I think that ambiguity aversion is not the sort of disposition we should want to instill in our robots, engineers and policy makers, risk aversion – another violation of classical expected utility maximization – may well be a feature of decision-makers that it is desirable for our robots, engineers and policy makers to display. So instead of showing that ambiguity averse agents, qua expected utility violators, can run in trouble, I will adopt a different strategy: my hope is to dispel the temptation to be ambiguity averse in the first place by suggesting that ambiguity aversion rests on a mistake in a way that other violations of expected utility theory, like risk aversion, do not.

In particular, I will suggest that ambiguity averse agents mistakenly avoid taking certain bets because they are treating their uncertainty about the probability of P as they would treat uncertainty about P, where in fact the two are quite different beasts that deserve very different sorts of treatment: When you're very uncertain about whether P (your credence is near 0.5), you should regard betting heavily on P as undesirable. But it's not true, I will argue, that when you're very uncertain about the *probability* of P, you should regard betting heavily on P as undesirable. Rather, when you're uncertain about the probability of P, you should be *uncertain* about whether or not betting heavily on P is desirable. Another way of putting the point: it's not the case that when probabilities are unknown it's best to be cautious. Rather, when probabilities are unknown, it's unclear whether it's best to be cautious.

Here's how I'll go about things: I've characterized ambiguity aversion relative to a set of probability functions *P* as the disposition to evaluate one's options in a way that gives extra weight to the probability functions in *P* that are pessimistic about that option. But if ambiguity aversion is a desirable feature of decision makers, which set of probability functions are we meant to be ambiguity averse relative to? I'll consider two interpretations: one on which each probability function in the set is a representation of some non-mental property that the world might (for all we know) have, and a second according to which each probability function in the set is, in some sense, an approximation of the agent's belief state.

3.1. The Objective Interpretation of Ambiguity Aversion

Here is the claim we'll consider first:

> OBJECTIVE AMBIGUITY AVERSION: In at least some contexts, it is desirable for decision makers to evaluate their options using a set of probability functions *O*, where each member of *O* represents a possible (non-mental) way the world might be (such as the ratio of marbles in the urn), and to evaluate each option in a way that gives extra weight to the probability functions in *O* that are pessimistic about that option.

Cases like the urn case may suggest that those who defend ambiguity aversion are thinking of the probability functions in question as objective probabilities. After all, the set of probability functions was generated by the possible ratios of marbles in the urn. But there are two ways that this set of probability functions could lead to ambiguity aversion. Here's the first way: the fact that we're uncertain about the ratio of marbles in the urn means that we, as agents, adopt a belief state that is *itself* best represented by a set of probability functions. Which set? The set including all and only the probability functions that correspond to the possible ratios of marbles in the urn.[5] What ultimately explains, on this picture, why we choose KNOWN, is the fact that we are ambiguity averse relative to a set of probability functions that represents our *belief state*. I'm going to discuss this way of thinking about ambiguity aversion under the subjective interpretation in the next section.

The second way that ignorance about which probability function represents some feature of the world might lead to ambiguity aversion is this: even if your own belief state is best represented by a single probability function, if you are ignorant of which member of a set of probability functions represents the relevant feature, you are nonetheless ambiguity averse relative to that set of functions. This is what I'll consider to be the genuinely objective interpretation.

Here is the problem with the genuinely objective interpretation: A central idea in decision theory is that which decision you should make depends on two things: (a) the values of the possible outcomes that can result from your decision and (b) your attitudes towards the propositions whose truth-value determines which outcome will come about. But the genuinely objective interpretation says that whether you should, for example, take a bet on P depends on more than that. On the objective interpretation, even holding fixed your attitude towards P (say you assign a 0.5 (subjective) credence to a red marble being drawn from UNKNOWN), and the values of the outcomes that will come about if you bet on P, whether you should bet on P depends on a third feature of your situation: whether you know some other fact Q (the ratio of red-to-black marbles in the urn). But this is very odd: how well you will do as a result of taking a bet on P depends only on *whether P*. This means that if your ignorance about Q is to make a difference to whether you decide to bet on P, it should do so by impacting your opinion about P itself. Once you've taken your ignorance about Q into account in forming your opinion about P, the mere fact that you're ignorant about Q should play no more role in a decision about whether to bet on P.

As an illustration of why this is so, consider the following decision theory: Whatever your opinion is about P, if you don't know what Jones thinks about P, don't bet on P. This is a silly theory. Presumably, the reason that ignorance of Jones' opinion might be relevant to whether you should bet on P is that ignorance of Jones' opinion about P is relevant to your own doxastic attitude

---

[5] While this is a common proposal in the literature on imprecise probabilities (see e.g. Joyce (2010), I believe this is the wrong way to generate the set that represents an agent's belief state. I explain why in Section 4.

towards P. But if your attitude towards P has already taken into account the fact that you don't know what Jones thinks about P, then, if it is nonetheless reasonable to be, say, highly confident that P, why shouldn't you bet on it?

The same holds for knowledge of objective probability: We should think of objective probabilities as a sort of expert, like your friend Jones. As with Jones, if you know that the objective probability of P is x, that will influence how you bet in virtue of influencing your belief state – by making *you* x-confident that P. And as with Jones, if you don't know the objective probability of P, that should influence how you bet in virtue of influencing your doxastic attitude towards P.[6]

What to make, though, of the fact that Aughenbaugh and Paredis purport to show that an objectively ambiguity averse decision procedure performs better, on average, than a decision procedure that is not objectively ambiguity averse? Before responding, let me explain how the experiment works (the details are a tad technical and readers who aren't interested in this experiment can skip to the next section). We're asked to imagine a team of engineers designing a pressure valve from a new type of steel. The yield strength of this type of steel is known to be normally distributed, but the mean and variance of the distribution are unknown to the engineers. The yield strength of the particular piece of steel that will be used for the valve is also unknown.

The engineers are now divided into two teams, which I'll call the "precisers" and the "imprecisers." Both teams are given 30 random samples from the true distribution (which has mean 180 MPa, variance 15 MPa). The precisers take their 30 samples and produce a single best-fit normal distribution. They then determine which design maximizes expected utility according to that distribution. The imprecisers use the data to generate a set of probability distributions corresponding to a 99% confidence interval. (Roughly, this means that if the experiment were performed a sufficiently large number of times, 99% of those times, the true distribution would be a member of the set). The imprecisers then use an ambiguity averse decision theory (Γ-maximin) over this set to choose a design.

Here's the computational procedure that Aughenbaugh and Paredis use to compare these approaches: First, they figure which design maximizes expected utility according to the true distribution. Call the expected utility of this design, *relative to the true distribution*, $EU_t(D_t)$. Second, they generate 30 random samples from the true distribution. Third, they figure out which design the precisers would choose upon observing these samples. Call this design D_p (for "precise design"). Fourth, they figure out which design the imprecisers would choose. Call this design D_i. Fifth, they calculate the expected utility of D_p and D_i using the *true* distribution. They repeat this procedure 100,000 times.

[6] One caveat: you might think that there are cases in which facts about objective probability impact the *value* of an outcome (see e.g. Bradley and Steffánson (2015)). The ways in which ignorance of objective probability might influence the value component of decision theory is an interesting topic, but not one I'll be exploring here.

The results are that the imprecise team tends to do better than the precise team. On average, the expected utility (according to the true distribution) of the design given by the precise team was 70, while the expected utility of the design given by the imprecise team was 89. This is a striking result which seems to lend strong support to the desirability of objective ambiguity aversion. But it is misleading. Let me explain why:

The reason that the imprecise team did better is attributable to three factors: (1) The precise team was not using the correct probability distribution. (2) The decision problem being worked with had the feature that the design that maximizes expected utility according to the true distribution is one that is conservative, in the sense that significant amount of money is spent to ensure safety. (3) Ambiguity averse decision theories tend to favor more conservative options.

Let's start with (1). The precise team shouldn't be using the best-fit normal distribution. Doing so amounts to ignoring the possibility that any of the other candidate distributions is the true distribution on the basis of a measly 30 samples. This is foolhardy. Instead, the precise team should be using what's called the *posterior predictive distribution*. This distribution is generated by modeling the team's ignorance over which mean and variance is correct using a uniform probability distribution and conditionalizing. The posterior predictive distribution generated from a small sample is going to be less opinionated than the best-fit normal distribution.[7]

This brings us to (2) and (3): the particular decision problem Aughenbaugh and Paredis are working with has the feature that a conservative design is the one that maximizes expected utility according to the true distribution. Objectively ambiguity averse decision theories will tend to favor conservative options. So if we choose a decision problem in which the design that maximizes expected utility according to the true distribution happens to be the conservative one, and we compare an ambiguity averse procedure with a precise overly confident procedure, it's no surprise that the ambiguity averse procedure will tend to do better.

To summarize: Aughenbaugh and Paredis' experiment does not support an objectively ambiguity averse decision theory. What it shows is that when a conservative design maximizes expected utility according to the true distribution, an objectively ambiguity averse decision procedure is more likely to yield that design when compared with an inappropriately confident precise procedure. In real life, we don't know which procedure will maximize expected utility according to the true distribution. If we knew that the conservative design were the preferred design, we wouldn't bother using *any* of these procedures. If we don't know which design is preferred according to the true distribution, Aughenbaugh and Paredis have given us no reason to

[7] Assuming the prior is uniform, the posterior predictive distribution will not be a normal distribution at all but rather a student-t distribution, which has fatter tails (displays more uncertainty) than a normal distribution.

think that the ambiguity averse procedure will yield better results than the (appropriate) precise procedure.

Let me end this section by responding to a potential objection to my treatment of Aughenbaugh and Paredis' experiment: Aughenbaugh and Paredis acknowledge that they don't account for the possibility of using a conditionalization-based procedure like the one I described. Their justification for this is that in many cases there are no grounds for choosing a prior distribution. Above I claimed that the precise team should be using a uniform prior. Aughenbaugh and Paredis would likely claim that this is choice is unmotivated – there is no basis on which to assign equal probability to the candidate distributions.

Response: if the thought is that assigning uniform probability over these distributions is problematic because certain distributions may deserve to be weighted more heavily than others, and the assignment of uniform probability ignores this fact, then ambiguity averse decision theories fare no better. For while the imprecise procedure does not explicitly assign equal probability to each distribution in the set, each of these distributions is being input into the decision theory in a way that assigns them equal weight. (It's true that *when evaluating a particular option*, ambiguity averse theories weight certain distributions more heavily than others, but this weighting is just a matter of the degree of to which the distributions are pessimistic with respect to the particular option – not a reflection of the epistemic desert of the function). Thus, the implicit assumption that none of the candidate distributions deserves to be taken more seriously than any other is embedded equally in both procedures.

In sum, the genuinely objective interpretation is committed to the view that whether we should bet on P depends on more than what we think about P – it also depends on whether we know some other fact, Q. This is a mistake, since how well we'll do with respect to a bet on P depends only on whether P. I argued that knowledge or ignorance about Q should impact whether we bet on P only by influencing our attitude towards P itself. I also argued that the Aughenbaugh and Paredis study, which purports to support objective ambiguity aversion, does not in fact do so. So let's move on to the subjective interpretation.

3.2. The Subjective Interpretation of Ambiguity Aversion

If someone tells me that a fair coin will be flipped, I will be 50% confident that it will land Heads. But when I consider propositions like: "The temperature high next New Year's in Boston will be less than 30 degrees Fahrenheit," it is implausible that there is a unique number between 0 and 1 that represents my degree of confidence in this proposition. As I mentioned in the introduction, a common strategy for dealing with such cases is to represent our degrees of confidence using a set of probability functions. Such a set is called an agent's "representor."

Which probability functions belong to the representor? Those that are, in some sense, precisifications of the agent's doxastic state. What does it take to be a precisification of the agent's doxastic state? Here is my favored view: p is a precisification of the agent's doxastic state if and only if, for all propositions A and B, if the agent is at least as confident in A as she is in B, then $p(A) \geq p(B)$, and if the agent is more confident in A than she is in B, then $p(A) > p(B)$.[8] (The notion of "at least as confident as" is taken as primitive).

Savage (1954) proves that if an agent's comparative confidence ordering is (a) complete[9], (b) satisfies certain plausible axioms (nontriviality, nonnegativity, additivity), and (c) there are "enough" propositions that the comparative confidence ordering is defined over[10], then a comparative confidence ordering gives rise to a unique probability function that is compatible (in the sense described above) with the agent's comparative confidence ordering. If, however, the agent has an incomplete ordering (and the relevant constraints are met), then there will be a set of probability functions that are precisifications of the agent: those probability functions that represent the possible completions of the agent's confidence ordering. This is the set I'll be thinking of as the agent's representor. [11]

The proposal I'll consider now is that the set of probability functions relative to which it is desirable to be ambiguity averse is the agent's representor.

> SUBJECTIVE AMBIGUITY AVERSION: If an agent has a set of probability functions S as her representor, then, when evaluating an option, it is desirable that she give extra weight to the probability functions in S that are pessimistic about that option.

---

[8] This notion of compatibility is what Fishburn (1986, p.336) calls "almost agreement." Note that it's compatible with almost agreement that, for some p in the representor, $p(A) > p(B)$ even though it's not the case that the agent is more confident in A than B. For other explications of what it takes to be a precisification, see Joyce (2010) and Rinard (2017).

[9] For all propositions A and B, either the agent is at least as confident in A as she is in B, or she is at least as confident in B as she is in A.

[10] More precisely, the ordering is "superfine": whenever the agent is more confident in A than in the null proposition, the set of worlds over which the comparative confidence ordering is defined, can be partitioned into propositions $B_1 \ldots B_m$, such that the agent is more confident in A than in B_i for all $B_i \in \{B_1 \ldots B_m\}$.

[11] Fishburn (1986) describes the Savage results and provides a nice survey of a number of closely related results. Joyce (2010) discusses similar results by Kraft et al. (1959) and Scott (1965) on p. 285. Joyce has raised a worry about understanding one's representor in terms of appealing to comparative confidence. The worry is that facts about the agent's judgments of stochastic independence won't be encoded by the set of probability functions. This problem can be solved by including *conditional* comparative confidence judgments in our set of judgments. Luce (1968) proves that a complete ordering of comparative conditional confidence judgments (that satisfies certain axioms, over an algebra of sufficiently many propositions), yields a unique probability function. If we weaken completeness, these judgments will yield a set of probability functions that encodes the agent's conditional and unconditional probability judgments.

Let H be the proposition that a fair coin will land Heads, and let L_{30} be the proposition that the temperature high next New Year's in Boston will be less than 30°F. Suppose I'm not (determinately) more confident in one than the other, that I have a precise credence in H, and an imprecise credence in L_{30}. If I'm given the chance to bet on H or bet on L_{30}, SUBJECTIVE AMBIGUITY AVERSION recommends betting on L_{30}.

Here's the problem with SUBJECTIVE AMBIGUITY AVERSION: if we're going to claim that there's something desirable about betting on H rather than L_{30}, we must have something to say that favors H over L_{30}. The ambiguity aversion proponent will say that what favors H over L_{30} is the fact we know the probability of H whereas we don't know the probability of L_{30}. But if we think about what "knowing the probability of H" amounts to on the subjective interpretation, we'll see that this provides no motivation whatsoever for preferring a bet on H.

To see why, consider the following analogy: If you're unsure which of Horse A or Horse B came in first, then you'll be unsure which of Horse A or Horse B it's better to bet on (assuming the winnings and costs of the bet are the same for both). *This is true even if you know exactly what time Horse A came in* (say it was 5:11). If, despite knowing what time Horse A came in, you're not sure whether Horse B came in before or after Horse A (perhaps you have no idea what time the race began) it would be a mistake to think that it's preferable to bet on Horse A on the basis of knowing the time that Horse A came in. For what matters is the *relative order* of Horse A and Horse B.

The same is true in the case of imprecise subjective probabilities: "knowing the probability of H" is like knowing the time at which Horse A arrived. On the comparative confidence picture of imprecise probabilities that I favor, the fact that we "know the probability of H" boils down to the fact that we have a variety of comparative confidence judgments between H and several other propositions – propositions that are completely unrelated to Boston weather. For example, we have comparative confidence judgments between H and propositions like "a coin weighted 0.51 towards Heads will land Heads" and "A coin weighted 0.49 towards Heads will land Heads." The fact that we can situate H within this ordering of propositions is analogous to our ability to situate Horse A's arrival within the ordering of minutes provided by the clock. Just like our ability to situate Horse A in the minutes-on-a-clock ordering doesn't motivate a preference for betting on Horse A, our ability to situate H in the coin-weighting propositions ordering doesn't motivate a preference for betting on H. On the subjective interpretation, there is nothing more to "knowing" the probability of a proposition than having a precise credence in that proposition, and there is nothing more to having a precise credence in a proposition than being able to situate it in a certain ordering. There is nothing bet-worthy about propositions that we can situate in certain comparative confidence orderings, and so there is nothing that tells in favor of betting on H over L_{30}.

Here is another way to bring out what's going wrong with subjectively-ambiguity-averse theories: Decision theories provide us with a way of evaluating the desirability of options. But goodness or value ultimately attaches to *outcomes*, not options. This means that any way of evaluating options must be answerable to our concern with what is actually valuable (well-being, world peace, chocolate bars or what have you) and must therefore satisfy certain constraints. I propose the following as one such constraint:

> Constraint: In a case in which there are only two outcomes, Good and Bad, *A* is a more desirable option than *B* if and only if the agent regards *A* as more likely to bring about Good (and hence less likely to bring about Bad) than *B*.

Here's the thought behind Constraint: Consider a case in which there are only two outcomes: Good, which is getting a teddy bear, and Bad, which is failing to get a teddy bear. In other words, we're imagining a case in which *the only thing* you care about is getting a teddy bear. In such a case, if you don't think *A* is any more likely to get you a teddy bear than *B*, then it's very mysterious in what sense your concern with getting a teddy bear could motivate a preference for *A* over *B*. And if your concern for getting a teddy bear (*the one and only thing* you value) can't motivate a preference for *A* over B, any decision theory that recommends *A* over *B* should be rejected.

Subjectively ambiguity averse theories violate Constraint. In URNS, there are only two possible outcomes: winning the prize and not winning the prize. By stipulation, you do not think the ambiguity averse recommendation (choosing KNOWN) is more likely to deliver the prize than a decision procedure which recommends indifference, or even one that recommends UNKNOWN. So it follows from Constraint that any decision theory which evaluates KNOWN as a better option than UNKNOWN should be rejected.

I think that what creates the illusion that it's preferable to bet on a proposition we're precise in rather than a proposition we're imprecise in is that when we lack a sharp probability in P, we may feel a sense of haziness and uncertainty similar to the sort of feeling we might have when we're uncertain about P itself. There are good reasons why when we're very uncertain about whether P (meaning our credence is in the 0.5 vicinity) it's better to not bet (heavily) on P. But there are no good reasons why, when we're uncertain about the *probability* of P (in other words, when we're imprecise), we're better off not betting (heavily) on P. A sense of uncertainty about the probabilities should lead to a sense of uncertainty about whether to bet, not to certainty in the desirability of not betting.[12]

[12] See Rinard (2015) and Mahtani (2018) for a defense of claims along these lines.

At this point you might think the following: I see why ambiguity would naturally lead to uncertainty about whether to bet, rather than certainty in the desirability of not betting. But maybe I have a tie-breaking rule for what to do when I'm unsure about whether to bet. A thought along these lines is proposed by Levi (1986, 122-140) and discussed by Steele (2007) and Bradley and Steele (2015). The idea is that when, at "the first level" neither option is favored, a tie-breaking consideration can be introduced: when there's a tie, take the more cautious option – the one that an ambiguity averse theory would recommend. Indeed, Bradley and Steele propose that ambiguity aversion may be valued in climate-policy decision making precisely because caution is thought to be of paradigm importance in these contexts.

The problem with this proposal is that the ambiguity averse option is not always the more cautious option. The case of Heads and L_{30} brings this out: if you lack a determinate confidence judgment between these two propositions, there's no sense in which betting on Heads is more cautious, just like, in the horse-racing case, there's no sense in which betting on Horse A is more cautious. For what is caution? Bradley and Steele say that cautious rules are those give special weight to the potential negative implications of a particular action (807). But this understanding of caution motivates *risk aversion* (more on that in a moment), not ambiguity aversion. The potential negative implications at stake whether you choose Heads or L_{30} are the same: failing to win the prize. So neither option is more cautious than the other. It's true, of course, that if you choose Heads rather than L_{30} you avoid an option that some probability functions regard as having a lower expected utility than the expected utility of Heads. But thinking that this makes Heads the more cautious option would be like thinking that the fact that some potential times Horse B came in are later than the time Horse A came in makes betting on Horse A the more cautious option. It doesn't. Betting on Horse A gives you no guarantee that you won't bet on the losing horse. It doesn't even make it more likely that you won't bet on the losing horse. In our case as well, neither of our options (Heads or L_{30}) provides any guarantee that a bad outcome (failing to win the prize) will be avoided, and neither option makes it any more likely that a bad outcome will be avoided. Caution, at least on any intuitive understanding of the word, does not, in general, motivate ambiguity aversion.

But wait – in the case of the Mayor, and in the design of the pressure valve, the ambiguity averse recommendation (not implementing the program, designing the more costly valve) *is* the more cautious one – the one that avoids the worst outcome. So maybe ambiguity aversion can be motivated by appeal to caution, at least in some cases? If so, is this sort of caution desirable?

I believe not. Insofar as it's desirable to be cautious, what we need is a decision theory that's *risk averse* – not a decision theory that's ambiguity averse. Risk averse decision theories are those that favor options that minimize the spread of *utility* across the different *possible outcomes*. Ambiguity averse theories, in contrast, favor options that minimize the spread of *expected utility* across the different *probability functions*. If we want to be cautious, we should be using decision theories that will assure avoidance of terrible outcomes – not decision theories that assure avoidance of expectedly terrible outcomes.

But why not be both risk averse and ambiguity averse? If it's good to be cautious, let's make sure to avoid terrible outcomes, *and* expectedly terrible outcomes!

This line of thinking is mistaken: if ambiguity aversion is motivated by appeal to caution, being both risk and ambiguity averse amounts to double counting. To see why, imagine you have a doctor that you've been seeing for years. The doctor knows you extremely well, she has an accurate sense of your attitudes towards risk, and you trust her completely. A certain blood test comes back positive and you're deliberating about whether to undergo a particular treatment. The treatment is usually safe and effective, but it has some risks associated with it. Precise probabilities concerning the risks are unavailable. Now, suppose you ask your trusted doctor for advice about what to do and she says: "I've looked through all the relevant medical journals and thought about your particular situation extremely carefully, but the evidence is complex and, as a result, I'm really not sure what to recommend." It would be a mistake to react to this by thinking: "Well, if the doctor isn't sure what to recommend, I better take the conservative option and forgo the treatment." This would involve double-counting your risk aversion. For the doctor's uncertainty about what to recommend was uncertainty *that already took your attitudes towards risk into account*. Because of her uncertainty about the treatment in question, she wasn't sure what to recommend, *given your risk attitudes*.

Similarly, when we have a set of probability functions and they make differing recommendations, then provided we have a decision theory (for single probability functions) that can account for risk, those risk attitudes have *already been accounted for in each individual probability/utility function pair's recommendation*. If the individual functions are *still* issuing differing recommendations, what that means is that your belief state is conflicted about what to do *taking your risk attitudes into account*. Weighing the probability functions in ways that issue conservative verdicts constitutes double counting.

4. Applications

I've argued that ambiguity aversion is not a feature of decision makers we should value. What are the implications of this for artificial intelligence, engineering, the demands of rationality and scientific reporting?

If our primary aim is to create robots that will make good decisions, or train engineers designing pressure valves, then the considerations above suggest that we have no reason to privilege the use of a set of probability functions rather than a single one.

Similarly, insofar as we, as agents, have a choice about which opinions to form, if our concern is making good decisions, there is no reason to try to maintain a state of ambiguity in our doxastic attitude if we find ourselves in such a state. We should expect to do no worse by "plumping" for a more precise state and acting in accordance with it.

Whether there are benefits to using imprecise probabilities in the case of scientific reporting is more complicated. One reason for this is that it's not always made explicit what the ranges of probabilities reported by scientists are meant to represent. Are they ranges of probabilities that are, in some sense, objective (they represent live hypotheses about some way the world might be), or do they represent the scientists' (subjective) estimate of the likelihood of some event?

Here's a way of seeing how these can come apart and why it matters: Suppose for simplicity that scientists have discovered and agreed upon the One True Model for some climactic system. This model will output the probability of some event E given a set of parameters. Let's call the probability that the model outputs given the actual value of all the parameters "the objective probability of E." And suppose there is one parameter that's unknown. Its value could be anywhere in [0,1]. Suppose also that, as it turns out, if the value of the parameter is v, then (given the other inputs), the probability of event E is v. This means that the possible objective probabilities of E are the numbers in [0,1]. Finally, suppose that some considerations favor higher values of v, some considerations favor lower values of v, and there is no obvious way to weight these considerations that yields a single probability distribution over the possible values of v. Thus, the scientists lack precise probabilities about the value of v, and, as a result, lack precise probabilities about E.

On the account of imprecise *subjective* probabilities I've been describing, it's extremely implausible that [0,1] represents *any* scientist's estimate that E will take place. For if 1 is in the representor, this means that the scientist lacks a comparative confidence judgment between E and the tautology. But surely, if a scientist claims to have no idea what the value of v is, it is a determinate fact about the scientist that she is less confident in E than she is in the tautology! It's also plausibly a determinate fact that the scientist is less confident in E than that a coin weighted 0.999 towards Heads will land Heads. Thus, any range that represents the scientists' estimate of the likelihood of E will be significantly narrower than [0,1].

What this means is that if all the policy makers are given is [0,1] – the range of possible *objective* probabilities – then they are not being given a set of probabilities that should be the input to a decision making process where the outcome of the decision depends on whether E (see section 3.1). This is a serious problem with a practice of reporting only the range of possible objective probabilities.

So if the scientists are going to report a range at all, we need them to report some range that represents their collective *subjective* estimate of the likelihood of E. This may seem like an impossible task. The scientists may claim that they simply have no idea what the value of the parameter in question is, or they may disagree with one another. So while it's true that they're definitely less than 0.99 confident that E will take place, are they less than .8 confident? 0.7? Any such number might seem arbitrary.

That there is no interval that obviously represents the scientists' collective estimate of the likelihood of E is, indeed, regrettable. But this should not prevent scientists from narrowing the range that policy makers will use. What we need scientists to provide policy makers with is an estimate that could rightly serve as the input to decision theory, and [0,1] is simply not that. So, uncomfortable as it may be, it is the responsibility of scientists as expert informants to do some narrowing, even if there is no particular narrowing that stands out as privileged.

What I've argued so far is that *if* scientists are reporting a range of probabilities, then that range must represent their subjective estimate of the likelihood of E, and so will frequently look, in some sense, arbitrary. But now the question is: if scientists are reporting something with these arbitrary features, is it better for that thing to be range, or a point in the range? I want to suggest that the status quo, which *discourages* (or even forbids) scientists from reporting precise probabilities in cases of ambiguity, may be sub-par. For if I'm right, when it comes time for policy makers to make a decision, we can expect them to do no better than act in accordance with a single probability function in the set the scientists have given them. And if somebody is going to be plumping for a single probability function, I hazard to guess that it would be better for it to be scientists than policy makers.

Some support for the claim that we're better off having the scientists plump than having the policy makers plump comes from a study done by Aidan Lyon (2017). Lyon's study shows that if you ask people to report a "best guess" probability (a single number), as well as a range of probabilities, the probability you get by aggregating the best guesses tends to be more accurate than the probability you get by aggregating the midpoints of the ranges. In other words, people's "best guess" probability contains more information than their range. This is some reason to think that if we required individual scientists to report a particular number constituting their best guess probability, and aggregated these guesses, decisions made on the basis of this procedure would be better than those made by policy makers who are just given a range of probabilities.[13]

[13] Another intriguing result from the Lyon study is that an aggregation procedure that did even better was one in which agents reported a point *and* a range. Lyon produced this aggregate by taking a weighted average of the points: the narrower the range accompanying the point, the more weight the point got. This result suggests that one thing people may be conveying by the width of their reported range is the amount or type of evidence their estimate is based on. (If people who have "better" evidence are reporting narrower ranges, weighting their estimates more heavily could result in a more accurate aggregate). If this result is robust, then it may support the usage of a *weighted* average of the scientists' point estimates. However, I hesitate to draw any firm conclusions at this point since I think it is, as of now, unclear exactly what people are reporting when asked to simply report a range of probabilities. I believe this is a fruitful area for future empirical work.

With all that said, it's worth noting that it may be valuable for policy makers to know the range of possible objective probabilities *in addition* to the scientists' subjective estimate. For the fact that there is a great deal of uncertainty about the objective probability of E will plausibly bear on questions concerning the value of gaining further information and the degree of resilience that the (subjective) probabilistic estimate should have. However, while it's valuable for policy makers to be aware that there is uncertainty about the objective probability, it is a mistake to assume that this information is best conveyed in the probabilistic estimate that will serve as the input for decision making. For, as I argued, the probabilistic estimate that should serve as the input for decision making should not be the range of possible objective probabilities.

A final caveat: there may be moral or political considerations that speak in favor of reporting exclusively ranges rather than points. These are important considerations that are beyond the scope of this project.

5. Conclusion

Ambiguity averse decision theories are those that, when evaluating an option, give extra weight to pessimistic probability functions amongst the relevant set. The purported desirability of ambiguity aversion has motivated the usage of imprecise probabilities in a number of disciplines. In this paper I've argued that the decision-theoretic motivation for imprecise probabilities is unsuccessful. The set of possible objective probabilities in a given decision-problem is not a set of functions that should serve as the input to decision theory. So if probabilistic ignorance should result in ambiguity aversion, this must go via the agent (or robot, or political body) adopting a belief state that is itself imprecise. However, once we understand what's involved in having an imprecise belief state, we see that there's nothing that privileges betting on a proposition we have a precise credence in rather a proposition one we're imprecise about. If we're going to continue to recommend imprecise probabilities in contexts in which decision making is important (engineering, artificial intelligence, policy-making), we need some account of what benefits we expect them to deliver. As of now, no good reasons have been provided that support the usage of a set of probability functions rather than a single one in decision making contexts.

References

Al-Najjar, Nabil and Weinstein, Jonathan. (2009). "The Ambiguity Aversion Literature: A Critical Assessment." *Economics and Philosophy* 25(3): 249-284.

Binmore, Ken. (2008). *Rational Decisions*. Princeton University Press

Bradley, Richard and Steele, Katie. (2015). "Making Climate Decisions." *Philosophical Compass* 10(11): 799-810.

Bradley, Richard and Steffánson, Orri. (2015). "How Valuable are Chances?" *Philosophy of Science* 82(4): 602-625.

Bradley, Richard, Hegelson, Casey and Hill, Brian. (2017). "Climate Change Assessments: Confidence, Probability, and Decision." *Philosophy of Science* 84(3): 500-522.

Buchak, Lara. (2010). "Instrumental Rationality, Epistemic Rationality and Evidence Gathering." *Philosophical Perspectives* 24(1): 85-120.

Buchak, Lara. (2013). *Risk and Rationality*. Oxford University Press.

Elga, Adam. (2010). "Subjective Probabilities Should be Sharp." *Philosophers Imprint* 10.

Ellsberg, Daniel. (1961). "Risk, Ambiguity and the Savage Axioms." *Quarterly Journal of Economics* 75(4): 643-669.

Fishburn, Peter. (1986). "The Axioms of Subjective Probability." *Statistical Science* 1(3): 335-358.

Gärdenfors, Peter and Sahlin, Nils-Eric (1982). "Unreliable Probabilities, Risk Taking and Decision Making." *Synthese* 53(3): 361-386.

Ghirardato, Paolo, Maccheroni, Fabio and Marinacci, Massimo. (2004). "Differentiating Ambiguity and Ambiguity Attitude." *Journal of Economic Theory* 118(2): 133-173.

Gilboa, Itzchak and Marinacci, Massimo. (2013). "Ambiguity and the Bayesian Paradigm." In *Advances in Economics and Econometrics: Tenth World Congress*, vol. 1, ed. D. Acemoglu, M. Arellano and E. Dekel. Cambridge University Press.

Gilboa, Itzchak and Schmeidler, David. (1989). "Maxmin Expected Utility with Non-unique Prior." *Journal Of Mathematical Economics* 18(2): 141-153.

Heal, Geoffrey and Millner, Antony. (2015). "Should climate policy account for ambiguity?" Centre for Climate Change Economics and Policy Working Paper No. 229, Grantham Research Institute on Climate Change and the Environment Working Paper No. 202.

Joyce, James. (2010). "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives* 24(1): 281-323.

Klibanoff, Peter, Marinacci, Massimo and Mukerji, Sujoy. (2005). "A Smooth Model of Decision Making Under Ambiguity." *Econometrica* 73(6): 1849-1892.

Kraft, Charles, Pratt, John and Seidenberg, A. "Intuitive Probability on Finite Sets." *Annals of Mathematical Statistics* 30(2): 408-419.

Levi, Isaac. (1986). *Hard Choices: Decision Making Under Unresolved Conflict*. Cambridge University Press.

Luce, R. Duncan (1968). "On the Numerical Representation of Qualitative Conditional Probability." *Annals of Mathematical Statistics* 39(2): 481-491.

Lyon, Aidan. (2017). "Vague Credence." *Synthese* 194 (10):3931-3954.

Mahtani, Anna (2018). "Imprecise Probabilities and Unstable Betting Behavior." *Noûs* 52 (1):69 87.

Mayo-Wilson, Connor. and Wheeler, Gregory. (2016). "Scoring Imprecise Credences: A Mildly Immodest Proposal." *Philosophy and Phenomenological Research* 93(1): 55-78.

Rinard, Susanna. (2015). "A Decision Theory for Imprecise Probabilities." *Philosophers' Imprint* 15.

Rinard, Susanna. (2017). "Imprecise Probability and Higher Order Vagueness." *Res Philospohica* 94 (2):257-273

Sahlin, Nils-Eric and Weirich, Paul. "Unsharp Sharpness." *Theoria* 80(1): 100-103.

Savage, Leonard Jimmie. (1954). *The Foundations of Statistics*. Wiley.

Schoenfield, Miriam. (2017). "The Accuracy and Rationality of Imprecise Credences." *Noûs* 51(4):667-685.

Scott, Dana. (1964). "Measurement Structures and Linear Inequalities." *Journal of Mathematical Psychology* 1(2): 233-247.

Seidenfeld, Teddy. (1988). "Decision Theory without "Independence" or Without "Ordering."" *Economics and Philosophy* 4(2): 267.

Seidenfeld, Teddy. (2004). "A Contrast Between Two Decision Rules for Use with (Convex) Sets of Probabilities: Γ-Maximin versuse E-admissibility." *Synthese* 140(1-2): 69-88.

Seidenfeld, Teddy, Schervish, Mark and Kadane, Joseph. (2012). "Forecasting with Imprecise Probabilities." *Internal Journal of Approximate Reasoning* 53(8): 1248-1261.

Steele, Katie. (2007). "Distinguishing Indeterminate Belief from "risk-averse" Preferences." *Synthese* 158(2): 189-205.

Steele, Katie. (2010). "What are the Minimal Requirements of Rational Choice? Arguments from Sequential-Decision Setting." *Theory and Decision* 68(4): 463-487.

Zaffalon, Marco and de Cooman, Gert. (2005). "Editorial: Imprecise probability perspectives on artificial intelligence." *Annals of Mathematics and Artificial Intelligence* 45(1-2): 1-4.