Can Imprecise Probabilities be Practically Motivated?

Abstract: The usage of imprecise probabilities has been advocated in many domains: A number of philosophers have argued that our belief states should be "imprecise" in response to certain sorts of evidence, and imprecise probabilities have been thought to play an important role in disciplines such as artificial intelligence, climate science, and engineering. In this paper I'm interested in the question of whether the usage of imprecise probabilities can be given a practical motivation (a motivation based on practical rather than epistemic, or alethic concerns). My aim is to challenge the central motivation for using imprecise probabilities in decision-making that has been offered in the literature: the idea that, in at least some contexts, it's desirable to be ambiguity averse. If I succeed, this will show that we need to reconsider whether there are good reasons to use imprecise probabilities in contexts in which making good decisions is what's of primary concern.

Word count: (excluding appendix, references and notes): 9997

1. Introduction

Suppose that Alice and Bob, who share all the relevant evidence, are contemplating whether neighbor Nadeem will wear a red shirt tomorrow. Alice adopts the the range of probabilities [0.1, 0.3] towards the proposition that Nadeem will wear a red shirt. Bob sticks his neck out and adopts the precise probability of 0.2. In this case, Alice adopts what are called "imprecise probabilities," whereas Bob's probability assignment is precise.[1] Question: should we expect Alice to do better than Bob in virtue of having adopted an imprecise rather than precise probability?

There are a number of ways in which an agent might be better off as a result of adopting certain probabilities over others: I'll mention just two: First, the agent might be more *accurate*: her belief state might better match, in some sense, the way the world really is. Second, the agent might

[1] Throughout, by "imprecise probabilities," I mean a (non-singleton) set of probability functions or a (non-singleton) set of numbers (typically an interval) in [0,1]. "Precise probabilities" are single probability functions or single numbers in [0,1]. By "adopting" a set of probabilities I mean entering a belief state represented by that set or taking that set as (part of) the basis on which one makes decisions. In the latter case, the formal upshot will be that the set of probabilities is one of the inputs to one's decision rule. An "imprecise agent" is an agent whose belief state is best represented by a (non-singleton) set of probability functions.

do better *practically*: she might be richer, happier, bring about world peace, or better achieve her practical aims, whatever they are. For each of these notions of "better" we can wonder whether it's preferable, for the purposes of doing "better," to be like Alice or like Bob.

A number of results suggest that accuracy considerations provide no reason to prefer imprecise over precise probabilities.[2] This paper concerns the question of whether the adoption of imprecise probabilities can be motivated by one's practical concerns.[3,4] Assuming that the world doesn't give out prizes directly for being imprecise, if it's sometimes better to adopt imprecise probabilities for practical reasons, it's because it's sometimes better to make decisions using imprecise probabilities rather than precise ones. And, indeed, imprecise probabilities have been thought to play an important role in decision-making in a wide range of disciplines spanning from engineering to artificial intelligence to climate science.

In engineering:

> "in engineering design decisions, it is valuable to explicitly represent the imprecision in the available characterization of uncertainties by using imprecise probabilities" (Aughenbaugh and Paredis (2005), p.1).

---

[2] Seidenfield et al. (2012), Mayo-Wilson and Wheeler (2016), Schoenfield (2017), Berger and Das (forthcoming), Builes et al. (forthcoming). For an opposing view, see Konek (forthcoming) who rejects some of the constraints on accuracy measures appealed to in these results.

[3] I don't intend to rule out the possibility that there can be motivations that are neither accuracy-based nor practically-based. For example, one may think imprecise probabilities are better reflections of the agent's evidence. See clarification 2 below.

[4] I want to refine the question a bit in order to set aside certain kinds of motivations for imprecise probabilities from my discussion. For example, an uninteresting way in which Alice might be better off practically by adopting imprecise probabilities is if it turns out that she gets great pleasure out of adopting these probabilities, or if God sends everyone who adopts imprecise probabilities to heaven. A more interesting way that sets of probabilities might be practically valuable, that I'll also be setting aside, stems from the idea that some of what we value might *itself* be probabilistic. For example, perhaps, as Stefánsson and Bradley (2015) and Bradley (2016) suggest, we value not just money, but the *chance* of getting a money. They propose a model on which sets of probabilities can play a role in determining the agent's *utility* function. While I think the idea that some of what we value is probabilistic in nature is an interesting one, when I talk about a "practical motivation" for using imprecise probabilities, I'm looking for a motivation for using sets of probabilities that doesn't depend on the particularities of an agent's utility function. In other words, the motivation I'm looking for would motivate using sets of probabilities in decision-making regardless of what sorts of utilities one assigns. Because I'm setting aside a motivation that's rooted in the idea that what we value is probabilistic, it will be easiest for exposition to just assume throughout that what we value is not probabilistic. This is unproblematic because if a motivation for using imprecise probabilities succeeds, in the sense I'm interested in, it should also motivate using imprecise probabilities in cases in which what we value is not probabilistic.

In artificial intelligence:

> [I]mprecise probability models are needed in inference problems characterized by scarce, vague or conflicting information…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:

Bradley, Hegelson and Hill (2017), in an article about the usage of imprecise probabilities in the report of 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 challenge the central motivation for the claim that adopting imprecise probabilities (rather than precise ones) is sometimes preferable from a practical perspective: a motivation stemming from what's called "ambiguity aversion."[5]

2. What Can Imprecise Probabilities Do for Us?

There's something intuitive about the thought that we shouldn't go around confidently making decisions on the basis of some particular probability function when, in some sense, the

---

[5] Let me take this opportunity to flag one other practical motivation for using sets of probability functions I'm aware of but that I won't be discussing here: Kyburg and Teng (1999) run simulations comparing the performance of a (precise) Bayesian method over a series of trials in a particular betting game they set up, and one based on the usage of confidence intervals from frequentist statistics. They reported results for six numbers of trials. When the number of trials was small (3, 5 and 10 trials), the precise Bayesian tended to be better, while with larger number of trials (20, 30 and 50) the confidence-interval method seemed to do better. (In the very long run they do equally well). This consideration won't clearly motivate using sets of probabilities in one-off cases (because for small numbers the precise Bayesian tends to do better), but it may motivate using sets of probabilities in certain circumstances where iterations are expected. Kyburg and Teng don't offer an explanation for this phenomenon, or an argument that it will hold in general. I myself am somewhat puzzled by it and I'd encourage those interested in comparing the performance of precise and imprecise methods to give this paper a careful look.

evidence leaves open several. But if we are to advocate imprecise probabilities in decision-making, we need to say something about what we expect the decision-maker to *do* with the set of probabilities. Sure, policy makers may not like to "base a policy decision on a flimsy information base" when a lot is at stake, but if they are to use a set of probabilities under such circumstances because scientists can't "fill in the gap," we need to consider how the policy makers are meant to act if the probability functions in the set issue differing recommendations. According to some decision rules that have been defended for imprecise probabilities, when it comes time to make a decision, an agent will simply act on the basis of one of the probability functions in the set.[6] But if at the end of the day we're going to "plump" for one and act on its basis, what's the point of lugging around the whole bunch?

What we need, then, to motivate the usage of imprecise probabilities in decision making is a theory on which the decision-maker who adopts sets of probabilities will act in a way that's (at least sometimes) recognizably different from, and that we regard as preferable to, the way a decision-maker operating with a single probability function would act. My focus in this paper will be the central class of decision rules which have been thought to fit this bill: ambiguity averse rules. (Indeed, in the literature on decision rules for imprecise agents, ambiguity averse rules are the *only* rules I'm aware of that make imprecise agents behave in a way that is both recognizably different from the behavior of precise agents *and* has been thought to be preferable from a practical perspective).

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 Alice, who uses the entire interval [0.1, 0.3] when making decisions, and that the extra caution that Alice will display is (in at least some contexts) desirable. (A more precise characterization of ambiguity aversion will be given in the next section).

Before proceeding, some clarifications about the project:

Clarification 1: If you're familiar with a certain strand in the literature on imprecise probabilities, which focuses on the potentially *negative* consequences of moseying around in an imprecise fashion,[7] the question of whether we can expect there to be any *benefits* from imprecise probabilities might strike you as odd. I'm going to set these types of worries about imprecise

[6] For example, Levi (1980).

[7] Seidenfeld (1998) and (2004), Al-Najjar and Weinstein (2009), Elga (2010) and Steele (2010).

probabilities aside. Maybe there are negative consequences for adopting imprecise probabilities under certain circumstances. But if the usage of imprecise probabilities 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 we might still ask: can we expect imprecise probabilities to have benefits that could, all things considered, motivate their usage?

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 do better practically 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 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 the decision-making apparatus?

Clarification 3: I'll be idealizing agents in the following respects: I'll assume that an agent's degrees of confidence (her "credences") can be represented by a probability function or a set of probability functions, that she updates by conditionalization, follows her decision-theory perfectly, and has no normative or evidential uncertainty (uncertainty about which probability function or functions are supported by evidence, what her evidence is, what decision-theory should be followed, or what the theory recommends in some particular case). These assumptions are of course idealizations – I'm not claiming that *we* should assume our decision-makers will be ideal in these ways, or even that they ought to be.[8] There may well be reasons to use sets of functions grounded in the idea that decision-makers won't satisfy these constraints. I don't know what they would be, but it may be an interesting avenue to consider.

Clarification 4: My aim is to challenge the claim that ambiguity aversion is a *desirable* feature of decision-makers, and thereby a motivation for thinking that it is desirable for imprecise probabilities to be used in contexts in which decision-making is what's of primary concern. There are a number of arguments in the paper, and while they are primarily intended to challenge the desirability claim, some may be thought to support the stronger conclusion that it is in fact

[8] One might think, a la Christensen (2008), that even perfectly rational agents will exhibit normative and evidential uncertainty.

undesirable or *irrational* to be ambiguity averse. My aim in this paper, however, is only to defend the weaker claim: ambiguity aversion isn't something to aspire to, and so there is no good reason, stemming from ambiguity aversion, to favor the usage of imprecise probabilities in decision-making. I will leave it the reader to consider which arguments, if any, support the claim that it is in fact *un*desirable to be ambiguity averse.

One last point: I'm using the term "desirable" rather than "rationally required" in part because I'm interested in what features we want systems that will be making decisions on our behalf to have. These systems can be individual agents, groups of agents, or machines. How or whether to attribute rationality to groups or machines is a vexed question that I'd prefer to skirt. So throughout the paper, when I ask whether ambiguity aversion is "desirable" you can read that as code for: should we prefer and thereby encourage policy-makers, engineers, robots, our future selves, or any other system that will be making decisions on our behalf to display ambiguity averse tendencies? Since ambiguity aversion is the central motivation that has been offered for preferring imprecise probabilities in decision-making contexts, if it turns out that the answer to this question is "no," we need to reconsider the (admittedly intuitive) thought that when the evidence doesn't pick out a single probability function, decision-makers should avoid making decisions on the basis of one.

3. Ambiguity Aversion

This section describes ambiguity aversion and its motivating examples –familiar readers may wish to skip it.

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 should go 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. Ambiguity aversion, very roughly, can be thought of as a preference for acting on known rather than unknown probabilities.

To get a feel for the phenomenon, it will be helpful to consider 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 prefer to draw from KNOWN.

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 50%. The mayor decides to go ahead with the program. If she were told that the probability of a wet winter were anything above 50%, 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, agents shy away from taking a gamble when the probabilities are "unknown" (what such ignorance amounts to is a topic we'll return to). However, a simple preference for taking gambles when probabilities are "known" rather than "unknown" won't quite capture the phenomenon that ambiguity aversion proponents are advocating: those who defend the desirability of ambiguity aversion will want to say that how an agent should act when the range of possible probabilities is 40-60% is different from how the agent should act when the range is 20-80% even though in both cases the probabilities are unknown. So to capture the phenomenon, we need to say something a bit more complicated: the preference in question is a preference for acts "whose evaluation is more robust to the possible variation of probabilities" ((Klibanoff et al (2005) p. 1852). Two questions: First, which probabilities? Second, what does a preference for acts "whose evaluation is more robust to the possible variation of probabilities" amount to and how does it capture the verdicts in the cases described above?

Let's start with the first: Which probabilities are the ones such that it's desirable (on this view) to be "robust" relative to? Answer: those that the agent is or ought to be in some sense "uncertain" between. So we'll say somewhat informally:

An agent is ambiguity averse if the agent prefers acts whose evaluation is more robust to the variations of probabilities she is "uncertain" between.[9]

Which ones are those? And what does it mean to be uncertain about a probability function? This will be a choice point that I'll talk about much more in what follows. But the general thought is that certain features of an agent's circumstance (like, for example, her evidence), pick out *some* set of probabilities that the agent is "uncertain" over and that it is desirable to choose acts whose evaluation is robust to the variation in that set. So, at a general level, we'll say that

An agent is ambiguity averse relative to a set of probability functions P if the agent prefers acts whose evaluation is more robust to the possible variation of probabilities in P.[10]

Next: what does a preference for acts whose evaluation is "more robust to the possible variation of probabilities in *P*" mean? Here's one way to capture the thought: consider a set of probability functions *P*. Suppose, just for the purposes of illustration, that each probability function evaluates an act by calculating its expected utility. Now let's order the probability functions on the basis of how much they approve of that act – how much expected utility they assign to it. At one

[9] This definition requires a small qualification: it is not intended to apply to agents who prefer such acts because of particular facts about their *utility* function. This qualification is important if ambiguity aversion is to motivate the adoption of imprecise probabilities in the sense I'm interested in. See note 4 for more detail.
[10] The qualification referred to in the previous note applies here as well.

end we'll have the pessimistic probability functions – those that assign the option low expected utility, and at the other end we'll have the optimists – those that assign the act high expected utility. If we now ask: how desirable is the act according to *P as a whole*, we'll need some way to take into account the fact that the members of *P* evaluate the act in different ways. If we consider the evaluation of the act according to the *most* pessimistic probability function, and that act ends up looking like pretty good, we know that it will look even better according to the optimistic probability functions in *P*. In this sense, evaluating the act the way the pessimist does leads to an evaluation that is robust across the range of probability functions. But we don't need to give the most pessimistic evaluator *all* the weight: we can think of the extent to which we give more weight to more pessimistic evaluations as the extent to which we favor acts that are robust across a range of probabilities.[11]

The most extreme ambiguity averse theory, Γ-maximin,[12] tells us to evaluate each act in accordance with the probability function in the set that is *most* pessimistic about that act: the probability function that assigns that act lowest expected utility. It then recommends that we choose the act that gets the best of these evaluations: the act 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. (Less extreme theories that allow pessimists to get extra weight, but not all the weight[13] will also deliver the ambiguity averse verdict in URNS).

Al-Najjar and Weinstein (2009) say that "a leading interpretation of the ambiguity aversion literature is that the [ambiguity averse] choices are *rational* responses by decision-makers to a lack

---

[11] There are a number of models of ambiguity aversion in the literature, and representation theorems have been proven with respect to these models. See e.g. Schmeidler (1982), Gärdenfors and Sahlin (1982), Schmeidler (1989), Gilboa and Schmeidler (1989), Klibanoff et al. (2005), Binmore (2008), Ghiradato et al. (2004) and Gilboa and Marinacci (2011).
[12] Ellsberg (1961), Gärdenfors and Sahlin (1982).
[13] Ghiradato et al. (2004), Klibanoff et al. (2005), Binmore (2008) and Gilboa and Marinacci (2011).

of reliable information that prevents them from forming beliefs with confidence" (252, emphasis in original).[14] Epstein and LeBreton (1993) say that "[ambiguity averse choices] seem sensible at a normative level since they correspond to an aversion to imprecise information" (4). Bradley, Hegelson and Hill (2017), in their discussion of ambiguity aversion, take the upshot of a large body of literature on the subject to be that "insistence on a single precise probability…may have unintuitive, and indeed normatively undesirable consequences" (504) and they describe the kinds of choices that ambiguity averse rules recommend as seeming "appropriate for important policy decisions" (508). In some of the engineering applications I described above, imprecise probabilities are advocated precisely because it is thought to be desirable to be ambiguity-averse. And, finally, Sahlin and Weirich (2014) who appeal to ambiguity averse rules in responding to Elga (2010) say that "pretending that information we have is better than it is rather than honestly taking the epistemic uncertainty into account is not very rational" (103).[15]

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.

My aim is to challenge the desirability of ambiguity aversion. In the course of doing so I will object to the idea that ambiguity aversion can be motivated by appeal to caution and, in the appendix, I will address the motivation offered by the Aughenbaugh and Paredis study. I won't, however, be engaging directly with case-based intuitions aimed at motivating the desirability of ambiguity aversion. So, for example, I won't engage directly with the intuition some people have that in cases like that of the mayor, implementing a policy might be reasonable when the probability is known to be 0.5 but unreasonable when all that's known is that the probability is in

[14] Al-Najjar and Weinstein actually argue *against* the rationality of ambiguity aversion, though for reasons that differ from mine. See note 18.

[15] Some of these authors are not totally clear about whether they are claiming that ambiguity aversion is desirable/rationally mandated, or merely permissible. As I say in clarification 4 above, my primary target is the desirability claim, since that's what's needed to motivate the idea that when evidence is equivocal it's *better* to use imprecise probabilities for decision-making purposes, or the claim that precise probabilities *shouldn't* be used in certain decision-making contexts. However, some of the considerations I raise may also tell against the permissibility claim.

[0.2, 0.8]. My view is that such intuitions are mistaken and my hope is that after reading this paper you'll agree. Nonetheless, if the intuitive force of the cases survives my arguments, then some adjudicating will need to be done between the more general theoretical considerations I appeal to and intuitions about cases. I will leave such adjudication to the reader.

4. Setting up the Challenge

There are different strategies one might deploy in raising a challenge to 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. For example, risk averse agents[16] are Dutch-bookable, will pay to avoid information, and will expect to do worse in the long run than a classical expected-utility maximizer.[17] Since ambiguity averse agents are not classical expected utility maximizers, they are also susceptible to such consequences.[18]

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. While I don't think that ambiguity aversion is 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 we want 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: the considerations I raise aim to dispel the temptation to be ambiguity averse in the first place. I will do this by raising a series of challenges to the idea that it's desirable to be ambiguity averse. I call them "challenges" because my primary aim is to convince you that at if we're going to encourage ambiguity aversion in decision-makers and decision-making

[16] Risk averse agents are those that have an all else equal preference for acts that minimize the spread of utility across possible outcomes.
[17] Buchak (2010) and Buchak (2013) Ch.7.
[18] Al-Najjar and Weinstein (2009) provide several examples to demonstrate this.

systems, several questions need to be answered that haven't been. I have no argument that the challenges I'll pose *can't* be answered. I raise them in the spirit of generating debate.

My starting point will be the following question: 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 an objective probability function that is consistent with the agent's evidence,[19] and a second according to which the set of probability functions is a representation of the agent's belief state.

5. The Objective Interpretation of Ambiguity Aversion

The section focuses on the claim that it's desirable to be ambiguity averse relative to a set of objective probability functions − probability functions that represent certain (non-mental) features of the world.[20] In what follows, I'll use "objective probabilities" synonymously with "objective chances" or just "chances." What exactly counts as "chance" is controversial. But for our purposes it's enough to point to some paradigmatic examples: the chance of a coin landing Heads is determined by how its weighted and the chance that a randomly selected marble from an urn will be red is the proportion of red marbles in the urn. If we were to be ambiguity averse relative to the set of candidate objective probabilities in URNS, we'd evaluate the act of drawing from UNKNOWN using the set of probabilities [0,1] − the set consisting of the possible chances of drawing a red marble.

Here's the claim we'll consider:

OBJECTIVE AMBIGUITY AVERSION:

(A) It is desirable for decision-makers to be ambiguity averse relative to a set of probability functions *O*, where *O* contains all and only the possible chance functions that are compatible with the agent's evidence.

---

[19] This interpretation fits well with the discussion in Klibanoff et al. (2005).

[20] Note that this very weak characterization of objective probability is compatible with thinking of what are called "evidential probabilities" as objective, at least if we think of evidential probabilities as just relations between propositions (e.g. Williamson (2000). If we were to allow for normative and evidential uncertainty one could have the following view: an agent's evidence always picks out a unique probability function, and an ideally rational agent's doxastic state will be representable by that function. Nonetheless even the ideally rational agent might be such that her evidence doesn't determine which of some non-singleton set of probability functions is the unique one supported by her evidence (see, e.g Christensen (2008) and Dorst (2020)). One might then claim that even though a rational agent's *belief state* will be precise, it is desirable that she be ambiguity averse relative to the set of candidate evidential probability functions that are compatible with her evidence. While I won't consider this view in detail (since it involves normative and evidential uncertainty, which I'm setting aside for the purposes of this paper) I do think a version of the worries I raise in this section would apply to it as well.

(B) The desirability of ambiguity aversion relative to *O* holds regardless of whether *O* represents or should represent the agent's belief state. (In other words, even if the agent's belief state is or should be represented by a single probability function, it is desirable that she be ambiguity averse relative to *O*).

The view that an agent's *belief state* should be represented by *O* and that it is desirable to be ambiguity averse relative to one's belief state, falls under the subjective interpretation discussed in the next section.

The challenge for OBJECTIVE AMBIGUITY AVERSION is as follows: It's widely accepted that we should treat chance like an expert. Here, for example, is Ned Hall: (1994, p.511): "…chance is like an expert in whose opinions about the world we have complete confidence." Indeed, chance has been flattered with all manner of expertise terms including "advisor" (Handfield, 2012), "guru," (Hall 2004), "guide to life," (Butler 1736, Lewis 1980), and "oracle" (Hall, 1994). But OBJECTIVE AMBIGUITY AVERSION recommends treating chance in a way that's very different from how we treat other kinds of expert opinion.

Here's why: The standard way we deal with the fact that we don't know what an expert thinks about P, in making a decision whose outcome depends on whether P, is to account for that ignorance when forming our credence in P (be it precise or imprecise) and then acting in accordance with it. OBJECTIVE AMBIGUITY AVERSION however tells us that when we're ignorant of a certain kind of expert's opinion about P, *chance*, we should not act on the basis of our credence in P, even if our credence is rational and has accounted for our ignorance of the chances. For example, OBJECTIVE AMBIGUITY AVERSION says that if you assign a 0.5 credence to a red marble being drawn from UNKNOWN, then, even supposing this credence is rational, you shouldn't take a bet at even odds that a red marble will be drawn. Despite the fact that your *credence* permits the bet, you should avoid it, because the pessimistic *chance functions* recommend against it. So, although we generally think it's fine to take bets on the basis of our credences when we don't know an expert's opinions, according to OBJECTIVE AMBIGUITY AVERSION, it is best to avoid bets when we don't know the chances.

It's worth emphasizing that we wouldn't *want* a decision-theory which told us, in general, to avoid bets when we don't know an expert's opinion. For bets don't just happen at casinos: taking an umbrella and eating a sandwich are bets too. Consider eating a sandwich. There is a possibility that eating the sandwich will make me sick. There are all sorts of experts whose probabilities I'd defer to concerning the proposition that the sandwich will make me sick. Still, the fact that I don't know the sandwich expert's opinion doesn't mean I shouldn't use my credences to decide whether to eat the sandwich. Rather, I should form my credences in a way that appropriately takes into account my lack of expertise, and whatever information I have, and then be guided by them. Note that any alternative way to account for ignorance of expert opinion would be quite a mess: often, when I make a decision, I'll know what some experts think, but I may not know what all the experts think, or what the best expert thinks. On the traditional picture, I take all the information I have about various expert opinions into account in setting my credences and use them to evaluate my options. But if we have "add-on" principles which say "don't use your credence in P when you're ignorant of an expert opinion about P" we'll face all sorts of questions about which expert opinions we need, what degree of expertise is required and so forth.

Another way to illustrate the fact that we don't want to circumvent our credences when we're ignorant of an expert's opinion arises when we think about one of the best experts around: the omniscient credence function which assigns 1 to all truths and 0 to all falsehoods (call it "God"). Consider the proposition that the randomly drawn marble from KNOWN is red. The credences that God might have in this proposition that are compatible with my evidence are {0,1}. If we should avoid bets when we don't know expert opinions, then I should avoid betting on a red marble being drawn from KNOWN because, while I know the chance, I don't know God's credence. Once God enters the pool of experts, thinking that we should avoid betting when we don't know an expert's opinion leads to the absurd consequence that whenever we're not certain of the truth value of a proposition we should avoid betting on it. This would result in a bland life indeed.

Summing up: OBJECTIVE AMBIGUITY AVERSION claims that a certain kind of expert – chance – is such that, ignorance of its opinion should guide our behavior by circumventing our credences. This, as I illustrated above, is not the way we do or should account for ignorance of expert opinion in general. So unless more is said about why ignorance about chance in particular should result in our evaluating a bet on P with probabilities other than our credence in P, OBJECTIVE AMBIGUITY AVERSION is treating ignorance about chance in a way that is *ad hoc*.[21]

[21] 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? The details are bit technical and not crucial for the main thread so I explain them in an appendix.

5. 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 (call this proposition "HEADS"). But when I consider propositions like: "There were between three and five birds in the garden at 6am this morning" (call this "BIRDS") it's implausible that there is a number between 0 and 1 that represents my degree of confidence in this proposition. A common strategy for dealing with such cases is to represent our degrees of confidence using imprecise probabilities. On this model, a doxastic state is represented by a set of probability functions (the agent's "representor") and a doxastic attitude towards a proposition is represented by a set of numbers (typically an interval).[22]

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. Views according to which certain bodies of evidence support sets of functions, an agent's representor ought to be the set of functions, *S*, supported by her evidence,[23] and it is desirable be ambiguity averse relative to *S*, fall into this category.

> SUBJECTIVE AMBIGUITY AVERSION: It is desirable that a decision-maker be ambiguity averse relative to the set of probability functions that is her representor.[24]

Challenge #1: Why treat probability different from utility?

According to SUBJECTIVE AMBIGUITY AVERSION, assigning an imprecise probability to an option can make that option less betworthy. But nobody, as far as I know, has defended the view that assigning an imprecise *utility* to an option makes an option less betworthy. Given that we don't

---

[22] The imprecise probability assigned to P is the set of probabilities assigned to P by each member of the representor.
[23] For example, Levi (1985) and Joyce (2005). Note that a view that I won't be considering (and that as far as I know nobody has defended) consists of the following two claims: (a) Bodies of evidence sometimes support a (non-singleton) set of probability functions, and (b) Even though an agent in possession of this evidence is or should be representable by a single probability function, the agent ought to *act* on the basis of (and be ambiguity averse relative to) the set of probability functions that the evidence supports. The reason I'm not going to address this view is that it seems unmotivated: if the evidence supports a set of probability functions, and you ought to be acting on the basis of that set of probability functions, why not think that your belief state (which is typically thought to take input from evidence and output action) should be representable by that set of functions? Furthermore, on at least certain functionalist stories of what it takes to have credences to begin with, an agent with evidence that supports a set of probabilities and who is acting on the basis of that set just *will* have a belief state representable by a set of probabilities. While I don't address this view explicitly, the third challenge I raise for subjective ambiguity aversion applies to this view as well.
[24] One may wish to add the qualification that the representor be supported by the agent's evidence. Nothing essential will rest on whether this qualification is added.

tend to think that imprecise utilities make an option worse, why think imprecise probabilities make an option worse?

To flesh out this challenge in more detail, it will be helpful to begin with a bit of background on imprecise utility: just like it's been thought that our doxastic attitudes can't always be modeled by a single probability function, it's been thought that our evaluative attitudes can't always be modeled by a single utility function. Ruth Chang's (2002) tea-and-coffee example illustrates the point: an agent might lack a preference between black coffee (call that "C") and tea (T). She might prefer coffee-with-sugar (C+) to black coffee (C), and yet also lack a preference between coffee-with-sugar (C+) and tea (T). Chang makes a point of saying that this can be so even if the agent knows all the relevant facts about the beverages: what they taste like, how energized she'll be after drinking them, etc.[25] There is no utility function that can represent this preference ordering. Just as sets of probability functions have been used to model agents whose beliefs aren't representable by a single probability function, sets of utility functions have been used to model agents whose preferences can't be modeled by a single utility function.[26]

Now consider the following: Suppose, just for the sake of simplicity, that utility varies linearly with dollars. You're at a coffee shop, trying to decide whether to buy a cup of coffee or save the two dollars. You've had the coffee many times. And let's suppose (perhaps unrealistically) that you know exactly what it will taste like, how you'll feel after you drink it, etc. Still, you don't assign precise utility to the option in which you buy the coffee (your utilities are imprecise, as discussed in the previous paragraph). Nobody would think to themselves: "Well, I know exactly what the utility of two dollars is, and I don't know what the utility of a cup of coffee is, so I'm better off saving the two dollars than buying the coffee." This would be a very odd way to think about your situation. In this case, failure to assign precise utility to the cup of coffee means you're unsure how to *compare* it with two dollars. If this is the case, you should be unsure about whether to buy the coffee or save the two dollars – not sure that you *shouldn't* buy the coffee. You might think that a similar line of reasoning applies to the choice between betting on BIRDS and HEADS. If you're

[25] In discussion of such an example Chang writes: "In this case, it is plausible to suppose that you know everything that is relevant to comparing the drinks and that in this case you have first-person authority over which tastes better to you…" (669)).
[26] See e.g. Levi (1986), Broome (1991, 92-3), Chang (2002) and Ok(2002).

unsure how to compare the probability of BIRDS and HEADS, you should be unsure about which to bet on – not sure that you *shouldn't* bet on BIRDS. [27]

Somebody may object to the analogy I'm drawing by pointing to ways in which probability and utility are different. Probabilities are bounded, utilities usually are not. Probabilities are revised through conditionalization in response to evidence while utilities are not. Prior probabilities reflect our opinions about the world and the ways of responding to evidence that we endorse, while utilities tend to reflect tastes and values. Some people think there is a unique rational prior probability function,[28] whereas very few defend such a view in the case of utility.[29] Any of these differences could potentially be a springboard for explaining why imprecise probabilities make an option less betworthy while imprecise utilities do not. I don't see at the moment why any of these features would motivate treating imprecision differently in the two cases, but I certainly don't want to claim that such a motivation couldn't be provided. So in raising this challenge, I'm inviting the defender of ambiguity aversion to offer such an explanation: to explain why whatever difference between probability and utility they think is relevant motivates thinking that imprecise probabilities make an option worse, but imprecise utilities don't.

Challenge #2: What does "not knowing" the probability of P amount to on the subjective interpretation and why is such lack of "knowledge" relevant?

Another challenge to SUBJECTIVE AMBIGUITY AVERSION questions the significance of probabilities being "known" on the subjective interpretation. The imprecise probability model as stated so far doesn't tell us anything about what mental states are involved in having precise versus imprecise probabilities. It just tells us that having. imprecise probabilities involves being in some state that is best represented by a set of probability functions – a representor. But which probability functions belong to the representor and why? My favored view is a comparativist one, which takes as fundamental an agent's comparative confidence judgments.[30] On this picture, an agent's

---

[27] See Rinard (2015) and Mahtani (2018) for a defense of claims along these lines.
[28] I do not endorse this view myself. But some who do include White (2005), Matheson (2011), Horowitz (2014), Greco and Hedden (2016), Horowitz and Dogramaci (2016), and Horowitz (forthcoming). For a survey of some of the relevant literature see Kopec and Titelbaum (2016)
[29] For a possible exception see Greco and Hedden (2016).
[30] This particular challenge won't be compelling to those who reject comparativism. (The other challenges I raise don't rely on this picture). For defenses of comparativism see Zynda (2000) and Steffánson (2017, 2018). An excellent survey by Konek (2019) discusses several versions of comparativism, representation theorems, extensions to the imprecise case as well as objections and responses to them. For an appeal to comparativist accounts of imprecise probabilities in particular see Schoenfield (forthcoming) and Builes et al. (forthcoming).

representor contains all and only the probability functions *p* such that: 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 $p(A) > p(B)$.[31] (The notion of "at least as confident as" is taken as primitive). When an agent's comparative confidence ordering is incomplete, this set of probability functions will contain more than one member. If it is complete, and certain other conditions are met, it will contain only one. So having imprecise credences, on this view, amounts to having an incomplete comparative confidence ordering.

If having imprecise credences just amounts to having an incomplete comparative confidence ordering, it's worth asking: what about such an ordering makes it the case that some of the propositions nonetheless get assigned precise probabilities (or have "known" probabilities) while others do not? Answer: the propositions that get assigned precise probabilities are just those that can be situated within some *sub*-ordering (an ordering of some subset of propositions in your algebra) that satisfies certain constraints.[32] So, for example, what makes it the case that you have a precise credence in HEADS is that you can situate HEADS in an ordering with certain other propositions – like the proposition that a coin weighted .49 towards Heads lands Heads, the proposition that a coin weighted 0.51 towards Heads lands Heads etc.

So here's what our question boils down to: Suppose you have a choice between betting on A and betting on B and you lack a comparative confidence judgment between A and B – you're no more confident in A than B or vice versa, and you're not equally confident in them either. Does the fact that you can situate A within a comparative confidence ordering of propositions: $\{C_1 \ldots C_n\}$, when you have no idea how to situate B in that ordering relative to A, make it preferable to bet on A than B? It's hard to see why it would. If you could anchor A in some ordering, and you *also* had some views about where B falls in that ordering, then situating A may enable you to make comparisons between bets on A and B that you couldn't make otherwise. But why would anchoring A, which you have no idea how to compare to B, in some ordering of propositions that

[31] This notion of compatibility is what Fishburn (1986, p.336) calls "almost agreement." For other explications of what is required for a probability function to belong to a representor see Joyce (2010) and Rinard (2017).

[32] In particular it satisfies the axioms required for deriving a unique probability function from an ordering. See Fishburn (1986) as well note 30 for further references.

you *also* have no idea how to compare to B, help you choose between betting on A and betting on B?[33]

On the comparativist interpretation, thinking that when you lack a comparative confidence judgment between HEADS and BIRDS, it's better to bet on HEADS than BIRDS because you "know" the probability of one and not the other is a lot like thinking that, when you can't tell which of Horse A or Horse B came in first, it's better to bet on Horse A than Horse B because you know what time Horse A came in on. Knowing that Horse A came in at 1:52pm, means you can situate Horse A's arrival within a certain ordering of events: the ordering of events provided by minutes of the clock. But assuming you have no views about when B came in relative to 1:52pm, the fact that you can situate Horse A's arrival within the minutes-on-a-clock ordering does nothing to favor betting on Horse A over Horse B. Similarly, if we think of "knowing" the probability of A as amounting to no more than having comparative confidence judgments that allow us to situate A in an ordering of some other propositions $\{C_1 \ldots C_n\}$, and we have no idea where B falls relative to A in that ordering, it's hard to see why such "knowledge" has any bearing on a choice between A and B.

I suspect that if SUBJECTIVE AMBIGUITY AVERSION remains tempting, even upon accepting a comparativist picture of what assigning imprecise probability amounts to, it's because we're thinking that if we can situate A in some ordering $\{C_1 \ldots C_n\}$ that somehow makes A the safer bet, or the more cautious bet, *even if* we don't have a view about where B is situated relative to A in that ordering. I think this is an illusion, just like it would be an illusion to think that betting on Horse A is safer than betting on Horse B when you can situate A in the ordering of minutes on the clock. I'll discuss the connection between caution and ambiguity aversion in detail towards the end of the following subsection.

Challenge #3: *What is a decision-maker's imprecision towards P an indication of and why does what it indicates devalue a bet on P?*

Recall that the question we're asking is this: Should we want robots, engineers, policy-makers, our future-selves or whoever will be making decisions on our behalf to display subjectively ambiguity averse tendencies?

---

[33] More carefully: if you're imprecise in B and precise in A, what will distinguish A and B is that for some $\{C_i \ldots C_j\}$ within $\{C_1 \ldots C_n\}$, you lack comparative confidence judgments between B and all the members of $\{C_i \ldots C_j\}$, while you have comparative confidence judgments between A and all the members of $\{C_i \ldots C_j\}$.

When we're asking this question, we're asking, more specifically, how we want our decision-makers to behave as a function of their credences and utilities. Let's focus on some particular credal state *c*. Holding utilities fixed, the question becomes: what are our preferences over acts *conditional on our decision-maker being in c*? Just to make things as simple as possible, we'll suppose the decision-maker's prior credences and utilities are the same as ours and that they have at least as much evidence as we do. We'll also assume that they will update on any additional evidence they receive by conditionalizing on it.[34]

Note that given background assumptions like the ones I just described, conditional on a more-informed decision-maker having a high credence in P, *we* have a high credence in P. So it's easy to see how conditioning on the *magnitude* of the decision-maker's credence (how high or low it is) will impact how we value a bet on P. As we'll see, though, it's not in general the case that conditional on our decision-maker being imprecise in P *we* are imprecise in P. It's also not true in general that we regard our decision-maker's imprecision concerning P as evidence *against* P. So the question at hand is this: what does conditioning on our decision-maker's imprecision do to our doxastic state which leads to our devaluing certain bets? Or, put another way: what are we taking a decision-maker's imprecision to *indicate*, which leads us to value a bet on P less highly than we would if she were precise?

The answer in the spirit of ambiguity aversion I take it is something like this: "More imprecise credences indicate a paucity of information." In other words, conditional on the decision-maker being imprecise, we have a high credence that the decision-maker doesn't have much information to go on. How does the lack of information devalue a bet? Perhaps as follows: "That there's less information available to the decision-maker means we'll want her to proceed cautiously. The fact that caution (in at least some circumstances) is desirable motivates the desirability of her acting in subjectively-ambiguity-averse ways."

My purpose in this section is to debunk each of the three underlined sentences, and so cast doubt on this line of reasoning.

[34] If their credences are imprecise they'll update their representor by conditionalizing each member of their representor on their evidence.

Response to: "More imprecise credences indicate a paucity of information."

While this thought may have some initial intuitive force, a number of examples that have been offered in the literature[35] show that there is no systematic positive correlation between imprecise credences and lack of information. In fact, it's often the reverse: imprecise credences are often had (and thought to be warranted) when there's more rather than less information. Here are some examples:

- At the beginning of a poker game you assign a precise credence to the proposition that Lydia has a straight flush on the basis of the ratio of straight-flush-hands to possible total hands. You then take note of Lydia's facial expressions, tone of voice, and the frequency at which she sips her beverage. Your credence that she has a straight flush becomes imprecise. (based on Weatherson (2002)).
- You're a detective investigating a crime. All you know when you walk into the office is that one of two suspects: entitled suspect A and suspect B committed a crime. At this point your credence is 0.5 that A committed the crime. You're then handed two large boxes: one has all the evidence in support of A's guilt, and another containing all the evidence in support of B's guilt: each box is full to the brim with documents, photographs, and lab reports. After examining all this evidence, your credence that each of A and B committed the crime spreads to [0.4, 0.6]. (based on Schoenfield (2012))

Note that it's not some idiosyncratic feature of these cases that credences become imprecise when there is more evidence. Often as we get more information, we get more *types* of information, and we also increase the likelihood that the information we get will be *conflicting*. These are features that often generate imprecision in one's credence. But even if you're not convinced that there's a positive correlation between *more* information and more imprecision, no case has, as far as I know, been made that imprecise subjective credences will systematically correlate with *less* rather than more evidence. At best we can say that more evidence will lead to more precision in some cases but less in others.

[35] Weatherson (2002), Schoenfield (2012) and Peden (2018)

Response to: "That there's less information available to the decision-maker means we'll want her to proceed cautiously."

This claim, I want to argue is false as well – or at least it doesn't hold in general. Even if you think imprecise credences *are* an indicator of something in the vicinity of less evidence (maybe it's not less evidence but less "weighty" evidence or more "ambiguous" evidence[36]), and even if *you* don't like making bets when you don't have much information, you shouldn't think that it would be best for your *decision-maker* to proceed cautiously when *she* has less information. One way to see this is as follows. Let H be the proposition: the decision-maker has a huge amount of information (or very "weighty" evidence), and let L be the proposition that the decision-maker has only a little bit of information (or very "non-weighty" evidence). Now, consider some proposition P, like the proposition that it will rain tomorrow, and let E be your total body of evidence at the moment. Let's compare your preferences conditional on H&E and conditional on L&E. Note that what's being conditioned on in both cases is, intuitively, the same amount (and weight) of information: E, and some proposition about the quantity (or weight) of the decision-maker's information in the future. So nothing about the idea that *you* devalue bets when the body of information you have or are conditioning on is small or non-weighty, motivates the thought that you should value a bet more highly conditional on H&E than conditional on L&E.

The more general point is that while conditioning on facts about the magnitude of a more informed agent's credence impacts the magnitude of your credence in a pretty straightforward way, conditioning on facts about the precision of a more informed agent's credence doesn't impact the precision of your credence in a straightforward way at all. Take the detective case described above. Suppose you're 0.5 that A committed the crime (you know nothing about either suspect). Conditional on your decision-maker looking at all the evidence and becoming [0.4, 0.6] that A committed crime, what's your credence that A committed the crime? In the most natural way of spelling out the case it will still be 0.5. It's natural to think that no individual credence function in your representor will take the fact that the decision-maker adopts a [0.4, 0.6] credence in A's guilt as evidence either for or against A. On the standard interpretation of imprecise probabilities, if all the credence functions in your representor will assign a 0.5 credence to A's guilt, conditional on

[36] See e.g. Keynes (1921), Joyce (2005), and Peden (2018). These notions are tricky to define so I won't attempt to do so here. But the point I'm making applies regardless of how one spells them out.

the decision-maker adopting [0.4, 0.6], then *you* assign a 0.5 credence to A's guilt conditional on the decision-maker adopting [0.4, 0.6].

In general, supposing the decision-maker is imprecise in P, doesn't automatically make *you* imprecise in P, or imprecise to the same degree.[37] Indeed, it's hard to see why the sorts of features that are thought to correlate with a decision-maker's adoption of imprecise credences – be it conflicting evidence of different types, a small amount of information, or non-weighty evidence – are features that correlate with a state of affairs such that conditional on it obtaining *you* are imprecise: *you* don't have conflicting types of evidence conditional on the decision-maker having conflicting type of evidence, *you* don't have quantity of information Q, or weight of information W, conditional on the decision-maker having that quantity or weight of information.

If we want to tell our decision-makers or program our AI to be ambiguity averse, we need some account of what our decision-maker's imprecision towards P indicates, and why what it indicates devalues a bet on P. I already pointed out that there's no systematic positive correlation between imprecision and lack of information. There's also no systematic correlation between the precision of the decision-maker's credences, and the precision of our own credences conditional on the decision-maker's degree of precision. But I haven't argued that there is *nothing* imprecise credences might indicate that could devalue a bet, so maybe you think there is something that a decision-maker's imprecision is an indication of, and that that thing, whatever it is, warrants proceeding cautiously, at least in certain contexts. This is what I'll address next:

Response to: "The fact that caution (in at least some contexts) is desirable motivates the desirability of acting in ambiguity averse ways."

Bradley and Steele (2015) propose that ambiguity aversion may be valued in climate-policy decision making because caution is thought to be of paradigm importance in these contexts. Can appeal to the importance of caution in certain contexts explain why it's desirable to be ambiguity averse?

To answer this question, it will be helpful to have some definition of caution in mind. Since Bradley and Steele explicitly appeal to the value of caution in explaining why ambiguity aversion might be desirable in certain contexts, I will use their definition. Following them, I'll think of a

---

[37] For much more careful argumentation in defense of the claim that you won't in general be imprecise in P conditional on a more informed agent being imprecise in P, see White (2009), Joyce (2010) and Schoenfield (2012). This is related to the point that "Reflection" fails when applied to an imprecise credence in a proposition.

cautious decision rule as one which gives extra weight to "the potential negative implications of a choice or action" (806). Are ambiguity averse rules those that give extra weight to the potential negative implications of a choice or action? Not in general. The case of HEADS and BIRDS brings this out: if you lack a comparative confidence judgment between these two propositions, betting on HEADS is no more cautious in the relevant sense than betting on BIRDS. This is because the potential negative implications at stake whether you choose HEADS or BIRDS are the same: failing to win the bet. So giving extra weight to the negative implications of an action doesn't privilege either of the two bets.

The above example shows that ambiguity aversion will sometimes recommend avoiding bets even though caution (in the sense defined by Bradley and Steele) would not. Ambiguity aversion also has nothing to say against actions that are *incautious* in this sense. For example, ambiguity aversion has nothing to say against paying \$5 dollars for a ticket that will pay you \$10 dollars if a fair coin lands heads and \$0 otherwise. But if you were inclined *even a bit* to weigh more heavily the negative consequences of an act (losing your \$5 investment), you wouldn't want to take this gamble. So ambiguity averse recommendations both tell you to sometimes avoid acts even though they're no less cautious than the alternative, and fail to tell you to avoid acts that *are* more cautious than alternatives.

What's the lesson to be learned from this? If we're understanding caution as the tendency to weight more heavily the negative consequences of an act, ambiguity aversion is simply not the right tool to capture this tendency: *risk aversion* is. Risk averse theories are specifically designed to give more weight to negative consequences.

But wait – in the case of the mayor, and in the design of the pressure valve in the Aughenbaugh and Paredis study, the ambiguity averse recommendation (not implementing the program, designing the more costly valve) *is* intuitively the more cautious one – it's 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 it is not. Here's a way of thinking about the relation between risk aversion and ambiguity aversion: Risk averse decision theories help us avoid terrible outcomes. They're more concerned with the low utility that could come about as a result of an action than the high utility that might come about as a result of the very same action. Ambiguity averse theories, in contrast, aren't designed to help us avoid terrible outcomes – they're designed to help us avoid *expectedly* terrible outcomes. Instead of being especially concerned with the lower *utilities* of an act that certain *states of the world* give rise to, ambiguity averse theories are especially concerned with the lower *expected utilities* of an act that certain *probability functions* give rise to. If we want to avoid terrible outcomes (we're cautious, in Bradley and Steele's sense), we should be using decision theories that will help us avoid terrible outcomes – not decision theories that will help us avoid 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 – giving extra weight to the potential negative consequences of an action – 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 how much extra weight you give to the potential negative consequences of an action – how "cautious" you are, in Bradley and Steele's sense – and you trust her completely. A test comes back positive and you're deliberating about whether to undergo a particular treatment. The treatment is usually safe and effective, but it sometimes has detrimental side effects. Precise probabilities concerning these side effects 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, *I've taken the extent to which you give extra weight to the potential negative consequences of an act into account*, 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: "I value being cautious, so if the doctor isn't sure what to recommend, I better take the conservative option and forgo the treatment." This would involve double-counting. The doctor's uncertainty about what to recommend was uncertainty *that already took your cautious attitudes into account*. Because of her uncertainty about the treatment in question, she wasn't sure what to recommend, *given that you're cautious*.

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,[38] caution, in Bradley and Steele's sense, will have already been accounted for in each individual probability/utility function pair's recommendation. If the individual functions

[38] For example, Buchak (2013).

are *still* issuing differing recommendations, what that means is that your belief state is conflicted about what to do *taking your tendencies towards caution into account*. Weighing the probability functions in ways that issue conservative verdicts constitutes double counting.

Let me sum up the point. If caution involves giving more weight to the potential negative consequences of an act, as Bradley and Steele suggest, then ambiguity aversion can't be motivated by appeal to caution. Caution, in this sense, motivates risk aversion, and adding ambiguity aversion on top in order to avoid bad outcomes amounts to double counting. Your theory will sometimes give more weight to the negative consequences of an action than you do. This means that if the defender of ambiguity aversion is going to use caution to motivate the theory, some alternative notion of caution is necessary. This alternative notion of caution, I assume, will have something to do with wanting to guarantee some minimum amount of *expected* utility, rather than guaranteeing some minimum amount of utility. Ultimately, however, what we want to avoid is (for example) a climate disaster, not an expected-climate-disaster.[39] If what we ultimately care about is avoiding the bad and not the expectedly bad, more needs to be said about why caution, in this sense, is something we want decision-makers to be concerned with, once they've *already* accounted for avoidance of the *actual* bad through risk aversion. So here too, I invite the defender of ambiguity aversion who wants to appeal to caution to say more: in particular, to tell us what the notion of caution at stake is, if not the tendency to weight more heavily the negative consequences of an action, why this other sort of caution is desirable, and why ambiguity averse decision theories do a good job of capturing it.

This section raised a number challenges to SUBJECTIVE AMBIGUITY AVERSION. These challenges can be summarized as a series of questions: First, given that we don't think it's better to avoid an option when utilities are imprecise, why think it's better to avoid an option when probabilities are imprecise? Second, what psychological state is represented by imprecise probabilities? If at the fundamental level we have just an incomplete confidence ordering, and that's what imprecise probabilities represent, the idea that we "know" the probabilities of some propositions and are ignorant of others is a bit misleading. The most we can say is that some propositions we can situate in certain kinds of suborderings and others we can't. But if you can't compare A and B, how does situating A within some C's, that you *also* can't compare to B, favor

[39] I'm assuming here that what we value is not *itself* probabilistic. See note 4.

betting on A over B? Third, if we want agents or machines who will be making decisions on our behalf to be subjectively ambiguity averse we need to ask: what could the state of their imprecision indicate such that conditional on that state obtaining a bet on P is less valuable than otherwise? The *precision* of a more-informed agent's credence doesn't impact the precision of our conditional credence in any systematic way: certainly not in the way the magnitude of a more-informed agent's credence does. And finally, when ambiguity aversion is being motivated by appeal to caution, what notion of caution is at play? The ordinary notion of caution, and the one given by Bradley and Steele in their discussion of ambiguity aversion, is one on which caution is a tendency to choose actions which avoid the worst kinds of outcomes. But caution in this sense motivates risk aversion, not ambiguity aversion, and using both theories at once with the aim of avoiding the worst kinds of outcomes constitutes double counting.

6. Conclusion

I've challenged the idea that ambiguity aversion, whether understood relative to objective or subjective probabilities, is a desirable feature of decision-makers. 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 ambiguity aversion offers no reason to privilege the use of a set of probability functions rather than a single one. Since ambiguity aversion is the central motivation for using imprecise probabilities in decision-making, undermining this motivation also undermines the central reason for thinking robots and engineers should be using sets of probabilities to begin with. Similar considerations apply to us as agents. Insofar as we have a choice about which opinions to form, if our concern is making good decisions, this paper undermines the central motivation that has been offered for thinking it is preferable, for these purposes, to maintain a state of ambiguity rather than "plump" for a more precise state and act in accordance with it.

Appendix: The Aughenbaugh and Paredis Study

In this study, 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. 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 out 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.

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 support objective ambiguity aversion. But it's misleading. Let me explain why:

First, 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 prior probability distribution over all of the possible values and conditionalizing.[40] What the posterior will look like will depend on the

[40] 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. A natural choice under conditions of complete uncertainty is to use a uniform prior. Aughenbaugh and Paredis might claim that even a uniform probability distrubtion under complete ignorance is unjustified. But why? If the thought is that, for all we know, some distributions may deserve to be given more weight than others, this is a problem for the procedure Aughbaugh and Paredis propose as well. For on their proposal all the probability functions are input into the decision procedure in a way that assigns them equal weight. (This is true initially. When evaluating a *particular option*, pessimistic functions *with respect to that option* get more weight – but this has nothing to do with the epistemic desert of the probability function). This is a choice point – one could very well apply their procedure in a way which favors some distributions over others. But uniformity is what they choose. So the

prior, but it's worth noting that if the prior is uniform – a natural choice under conditions of total ignorance – the posterior predictive distribution generated from a small sample is going to be less opinionated than the best-fit normal distribution, and so less likely to recommend risky options.[41]

Second, 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. [42]

The Aughenbaugh and Paredis experiment thus 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. And when we don't know which procedure will maximize expected utility, Aughenbaugh and Paredis have given us no reason to think that the ambiguity averse procedure will yield better results than the (appropriate) precise procedure.

---

seeming "groundlessness" of assigning a uniform prior probability distribution doesn't tell in favor of the imprecise ambiguity averse procedure over the precise Bayesian procedure.

[41] In fact, it 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.

[42] Still, you might wonder: does the fact that in *some* cases ambiguity averse theories perform better offer us a motivation for being ambiguity averse? No. For exactly a parallel argument can be run which tells us to prefer ambiguity *seeking* decision theories (those that give extra weight to *optimistic* probability functions). Here's how the argument would go: consider the set of cases in which the act that maximizes expected value according to the true probability distribution is the risky act. We can show that ambiguity seeking theories will yield the right recommendation in such cases. (This would show up in simulations, but it could also be noted that since risky acts are ones with a high spread of possible value, risky acts are ones which *might* yield (relatively) high positive value in comparison with conservative acts. This means that being guided by the optimist's evaluation of any given act will skew things in favor of risk). So what we know is that when the true distribution recommends one act, ambiguity aversion does better, and when the true distribution recommends another act, ambiguity seeking does better. Decision theory is supposed to guide us when we *don't* know what the true distribution is, which means that an argument in favor of using some particular decision theory can't appeal to features of the true distribution (like the fact that it recommends a conservative act). You might at this point claim that ambiguity aversion still fares better: since the true distribution *might* recommends conservative action, in the interest of caution, when we don't know the true distribution, it's better to be ambiguity averse. I discuss the connection between caution and ambiguity aversion on p.23-26.

References

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

Berger, D. and Das, N. (forthcoming). "Accuracy and Credal Imprecision." *Noûs*.

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. (2016). "Ellsberg's Paradox and Value of Chances." *Economics and Philosophy* 32(2): 231-248.

Bradley, Richard, Hegelson, Casey and Hill, Brian. (2017). "Climate Change Assessments: Confidence, Probability, and Decision." *Philosophy of Science* 84(3): 500-522.
Broome, John. (1991). *Weighing Goods*. Wiley-Blackwell

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.

Builes, David, Horowitz, Sophie and Schoenfield, Miriam. (forthcoming). "Dilating and Contracting Arbitrarily." *Noûs*.

Butler, Joseph. (1736), *The Analogy of Religion, Natural and Revealed, to the Constitution and Course of Nature*. Knapton.

Chang, Ruth. (2002). "The Possibility of Parity. *Ethics* 112 (4):659-688.

Christensen, David. (2008). "Does Murphy's Law Apply in Epistemology" In T. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology Volume 2*. Oxford University Press.

Dogramaci, Sinan and Horowitz, Sophie (2016). "An Argument for Uniqueness about Evidential Support." *Philosophical Issues* 26 (1):130-147.

Dorst, Kevin (2020). "Evidence: A Guide for the Uncertain." *Philosophy and Phenomenological Research* 100(3): 586-632.

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.

Epstein, Larry, G. and Le Breton, Michel. (1993). "Dynamically Consistent Belief Must Be

Bayesian." *Journal of Economic Theory* 61(1): 1-22.

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.

Greco, Daniel and Hedden, Brian (2016). "Uniqueness and Metaepistemology." *Journal of Philosophy* 113(8): 365-395.

Horowitz, Sophie. (2014). "Immoderately Rational." *Philosophical Studies* 167(1): 41-56.

Horowitz, Sophie. (forthcoming). "The Truth Problem for Permissivism." *Journal of Philosophy*.

Joyce, James. (2005). "How Probobilities Reflect Evidence" *Philosophical Perspectives* 19: 153-178.

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.

Kopec, Matthew and Titelbaum, Michael G. (2016). "The Uniqueness Thesis." *Philosophical Compass* 11(4): 189-200.

Konek, Jason. (2019). Comparative Confidence. In *The Open Handbook of Formal Epistemology*, edited by R. Pettigrew and J. Weisberg, 267-348. PhilPapers Foundation.

Konek, Jason. (Forthcoming). "Epistemic Conservativity and Imprecise Credence." *Philosophy and Phenomenological Research*.

Keynes, John, M. (1921). *A Treatise on Probability*. Macmillan.

Kyburg, Henry E. and Teng, Choh Man. (1999) "Choosing Among Interpretation of Probability" in Proceedings of the 15th Annual Conference on Uncertainty in Artificial Intelligence (UAI-99): 359–365. Morgan Kaufmann Publishers

Levi, Isaac. (1980). *The Enterprise of Knowledge*. MIT Press.

Levi, Isaac. (1985). "Imprecision and Indeterminacy in Probability Judgment." *Philosophy of Science* 52: 390-409.

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

Lewis, David. (1980). "A Subjectivist's Guide to Objective Chance" in R. Jefferey (ed.) *Studies in Inductive Logic and Probability Volume II*: 263-293. University of California Press

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

Matheson, Jonathan. (2011). "The Case for Rational Uniqueness." *Logic and Episteme* 2(3): 359 - 373.

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

Ok, Efe. (2002). "Utility representation of an incomplete preference relation." *Journal of Economic Theory* 104: 429–449

Peden, William. (2018). "Imprecise Probability and the Measurement of Keynes' "Weight of Argument." *IfCoLog Journal of Logics and Their Applications* 5(4):677-708

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.

Schmeidler (1982). "Subjective Probability without Additivity." *Procceedings of the American Mathematical Society.*

Schmeidler, David (1989). "Subjective Probability and Expected Utility without Additivity." *Econometrica* 57(3): 571-587.

Schoenfield, Miriam (2012). "Chilling Out On Epistemic Rationality: A Defense of Imprecise

Credences." *Philosophical Studies* 158 (2):197-219.

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

Schoenfield, Miriam. (forthcoming). "Meditations on Beliefs Formed Arbitrarily." In T. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology Volume 7*. Oxford University Press.

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

Stefánsson, H. Orri. (2017). "What's Real in Probabilism?" *Australasian Journal of Philosophy* 95(3): 573-587.

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

Weatherson, Brian. (2002). "Keynes, Uncertainty and Interest Rates." *Cambridge Journal of Economics* 26(1): 47-62.

White, Roger (2005). "Epistemic Permissiveness" *Philosophical Perspectives* 19(1): 445-459

White, Roger (2009). "Evidential Symmetry and Mushy Credence." In T. Szabo Gendler & J. Hawthorne (eds.), *Oxford Studies in Epistemology*. Oxford University Press. pp. 161-186.

Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford University Press.

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.

Zynda, Lyle. (2000). "Representation Theorems and Realism about Degrees of Belief." *Philosophy Of Science* 67(1): 45-69.