A counterexample to three imprecise decision theories

Seamus Bradley*

1 Abstract

There is currently much discussion about how decision making should proceed when an agent's degrees of belief are imprecise; represented by a set of probability functions. I show that decision rules recently discussed by Sarah Moss, Susanna Rinard and Rohan Sud all suffer from the same defect: they all struggle to rationalise diachronic ambiguity aversion. Since ambiguity aversion is among the motivations for imprecise credence, this suggests that the search for an adequate imprecise decision rule is not yet over.

2 Introduction

Imprecise credences are a growth industry in philosophy. Many people seem to have become disillusioned with the standard Bayesian epistemology, with its unreasonably precise belief states (see, for example Joyce (2010); Bradley (2014)). Relaxing the requirement that credence is represented by a single probability function is a popular suggestion. Such theories are known as theories of "imprecise credence".

*Department of Philosophy, University of Leeds, S.C.Bradley@leeds.ac.uk

How agents with imprecise credences ought to make decisions is a difficult topic. There are many decision rules that have been proposed, but each appears to have some undesirable features. This suggests a problem for imprecise credence. But why should an apparently epistemological view like imprecise credence be judged on decision-theoretic grounds anyway? First, imprecise credence views are seen as alternatives to orthodox precise probabilistic view on epistemology, and one of the nice things about the Bayesian view is that it fits very well with the standard maximisation-of-expected-utility view of decision making. Any replacement for orthodox Bayesianism should at least make some attempt to also provide a theory of decision making. Second, one major criticism of imprecise credence views has come from a decision-theoretic angle, and response to this criticism requires discussion of decision making. Elga (2010) in fact argues that there is no adequate decision rule for agents with imprecise credences. This, he takes to be an argument against using imprecise credences at all. We will explore Elga's argument in the next section, and then discuss some decision theories that avoid the pitfalls of that example. These include the decision theories advocated by Moss (2015), Rinard (2015), and Sud (2014). We will then look at another decision problem, and show that those decision rules that do well in the Elga case will fare poorly in this case.

In the remainder of this section, I introduce the precise and imprecise credence models that are at issue in this paper. Both precise and imprecise epistemologies share the idea that the objects of belief for a rational agent – whom I shall call "you" – are elements of a Boolean algebra of propositions. That is, we have a set of states W – sometimes called a set of possible worlds – and propositions are subsets of that set. I shall assume throughout that W is finite.

The standard Bayesian epistemology has it that your credal state at a time is represented by a probability function, p, which is a non-negative, real valued function on an algebra of propositions which satisfies the following two properties:

- $p(W) = 1$
- If $x, y \subset W$ are disjoint then $p(x \cup y) = p(x) + p(y)$

Your degree of belief in a proposition x is normally identified with your probability of

x, that is, with $p(x)$.

A *gamble* is a function from a state space to the real numbers. The bigger a number the gamble assigns to a state, the better it is, if that state occurs. For example, if we're betting on the roll of a (fair, six-sided) die, and you bet that it will land on an even number (at odds of 1 to 1) then the state space is "Die lands 1, Die lands 2..." and the gamble is the function that returns 1 on the even numbers and -1 on the odd numbers. Throughout the discussion I make some standard assumptions: gambles are denominated in utilities, and you are neutral with respect to concerns of risk. The expected value of a gamble is the sum of probability weighted values of the states. That is: $E_p(f) = \sum_{w \in W} p(\{w\}) f(w)$.

Imprecise credence instead represents your credal state by a *set* of probability functions, P. Following van Fraassen (1990), we call this set P your *representor*. It is somewhat standard to understand your degree of belief in x as represented by the set of values assigned to x by members of P, which we write $P(x)$. That is, we interpret $P(-)$ as a set-valued function that outputs $P(x) = \{p(x), p \in P\}$. People often talk about this function as if it was a good representation of your imprecise belief, though arguably we should take the set of functions to be the genuine representing object.[1] Imprecise expected values are also determined pointwise: $E_P(f) = \{E_p(f), p \in P\}$.

3 The Hammond/Elga decision problem

Isaac Levi has long been an advocate of imprecise credences (Levi 1974; 1980). He advocated a decision rule known as "E-admissibility" whose particulars won't matter here. Peter Hammond suggested that E-admissibility had a problem with sequences of choices (Hammond 1988). Adam Elga has recently used a similar example to argue that *no* imprecise decision theory is acceptable (Elga 2010). It seems like an advocate

[1]We don't have the space to fully argue for this claim here, but consider a proposition about which you have no evidence, x. $P(x) = [0,1] = P(\neg x)$, but focus on the set of values misses the important fact that your attitudes to the propositions x and $\neg x$ are *complementary* in the sense that for every $p \in P$, $p(x) = 1 - p(\neg x)$.

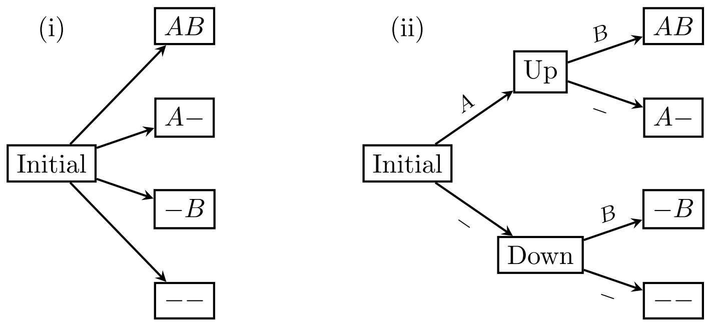

Figure 1: One-shot (i) and sequential (ii) versions of Elga's problem (after Bradley and Steele (2014))

of imprecise credences ought to have something to say about this decision problem.

The problem goes roughly like this. (This formulation is closer to Elga's version). You are about to be offered two bets on some unknown proposition h, one after the other. You know nothing at all about h: your credence in h is unconstrained. That is, $P(h) = [0,1]$. The bets pay out as follows:

- A – you lose 10 if h is true and win 15 otherwise
- B – you win 15 if h is true and lose 10 otherwise

To see the problem, let's start with a simple decision rule: maximise minimum expectation.

If we had a four way choice between "accept both bets", "accept A only", "accept B only", and "accept neither bet" – as in Figure 1(i) – then the imprecise expectations of the four options would be 5, $[-10,15]$, $[-10,15]$ and 0 respectively. Maximising minimum expectation – maximin – would recommend accepting both bets. So far so good.

But note that the decision problem above is not a four-way choice, but a sequence of two two-way choices (as in Figure 1(ii)). Let's walk through what goes wrong in the sequential choice. Let's start with maximin again. In a choice between accepting bet A or refusing it, a maximin-chooser would opt to refuse the bet: refusing has minimum expectation 0, versus -10 for accepting the bet. And likewise for bet B. So the agent ends up with refusing both bets. But this seems like a mistake: after all, the agent would be better off by having chosen to accept both bets. Thus, Elga concludes, the maximin decision rule is flawed. The point here is that this decision problem is a problem about *sequences* of decisions. The decision made at each node individually is not what is being criticised: it's how the decisions at different points in the choice tree *fit together*.

Sahlin and Weirich (2014) point out that there is another way of cashing out maximin – a more sophisticated way – that avoids this problem (see also Chandler (2014); Bradley and Steele (2014)). The basic insight is to use backwards induction to solve the decision tree. Consider the choice at node Up. You prefer accepting the bet (and thus accepting both bets, expectation 5) to refusing the bet (and thus having accepted bet A only, expectation $[-10, 15]$). And at Down you prefer refusing the bet B to accepting it (0 versus $[-10, 15]$). At the initial node, taking into account these facts about how you will choose in the future, you prefer accepting the first bet (and then accepting the second bet) to refusing the first bet (and then refusing the second bet). So the sophisticated approach to maximin choice doesn't have the flaw that the naive approach did. In an erratum to his paper on his website, Elga concedes that sophisticated maximin is immune to the criticism of Elga (2010). This sophisticated maximin approach, however, is unpalatable for other reasons that we shall discuss later on in the paper. So, the search is on for an alternative decision rule that can avoid this problem.

Let's now meet our new proposed decision theories, to see how they solve this decision problem. Inspired by this sophisticated maximin strategy, the forward looking rule (Sud 2014) assesses your options now, taking into account the sequences of future choices that they could be part of, and you act now with a view to avoiding those sequences you recognise to be inadmissible. A sequence of acts is inadmissible if it does not maximise expected utility for any $p \in P$. Since refusing bet A is part of an inadmissible sequence (refuse A, refuse B), while accepting bet A is not, this rule makes accepting bet A obligatory. The forward looking rule thus also avoids the Hammond/Elga problem, although it is still subject to the version of the problem outlined by Hammond (1988) (see footnote 18 of Sud (2014)). In outline, that problem involves a sequential decision where neither choice at the initial branch is such that it is not part of an inadmissible sequence.

Moss (2015) develops her identification theory with avoiding Elga's problem in mind. So it's not surprising that identification theory never sanctions refusing both bets. The key idea is that in sequences of choices, an agent "identifies with" a particular member of P and acts so as to maximise expected value with respect to that precise probability. You can change which p you identify with, but it is rare that you switch which p you identify with between being offered bet A and bet B, because by stipulation very little time elapses between the two bets. Moss argues that this is reasonable: that changing which p you identify with is something that happens rarely, and not in any sort of systematic or strategic way. No precise $p \in P$ makes refusing both bets permissible, so as long as you identify with the same p throughout the sequential choice, you are guaranteed to avoid the sure loss of refusing both bets.

Susanna Rinard's view starts from the perspective of considering a set of probabilities in a supervaluationist light. Something is determinately permissible only if it is permissible according to all $p \in P$. According to some $p \in P$, rejecting bet A is permissible, and according to some others, it is not permissible. So on the supervaluationist view of Rinard (2015), it is indeterminate whether it is permissible to reject bet A. And likewise for bet B. But, according to Rinard, it is determinately impermissible to reject both bets, since no $p \in P$ makes the "compound action" (refuse A, refuse B) permissible. So Rinard's rule also avoids Elga's problem.

We've seen three decision rules that, in different ways, make impermissible refusing both bets in the Hammond/Elga problem. In the next section, we introduce a new decision problem and show that it is precisely the same feature that helps these rules succeed in the Hammond/Elga problem that causes them to fail in cases of diachronic ambiguity aversion.

4 Rationalising ambiguity aversion

Many find it reasonable that you can prefer risky gambles to ambiguous gambles. A gamble is *risky* if the probabilities of the relevant states are known, whereas a gamble is *ambiguous* if the probabilities of the relevant states are unknown, or partially known.

Imprecise credences can tell a story about the rationality of this *ambiguity aversion*, and doing so seems to be a desideratum for imprecise credences.

The best known example of a decision problem that evinces this aversion for ambiguity is the "Ellsberg problem" (Ellsberg 1961). There is a somewhat intuitive and experimentally well-confirmed pattern of preferences that cannot easily be accommodated within the standard precise Bayesian framework (Fox and Tversky 1995; Camerer and Weber 1992).[2] This problem is sometimes used to motivate imprecise credences (e.g. Halpern (2003), p.24 ff.; Gärdenfors and Sahlin (1982); Levi (1986)). Here is a version of the problem.

> I have an urn that contains ninety marbles. Thirty marbles are red. The remainder are blue or yellow in some unknown proportion.

We're going to consider some bets that win 1 utility if the event in question occurs and nothing otherwise. Consider a choice between a bet that wins if the marble drawn is red (I), versus a bet that wins if the marble drawn is blue (II). You might prefer I to II since I involves *risk* while II involves *ambiguity*. Now consider a choice between a bet that wins if the marble drawn is not blue (III) versus a bet that wins if the marble drawn is not red (IV). Now it is III that is ambiguous, while IV is unambiguous but risky, and thus IV might seem better to you if you preferred risky to ambiguous prospects. Call a preference for risky bets in the two pairwise choices the *Ellsberg preferences* or Ellsberg choices. Such a pattern of (strict) preferences (I preferred to II but IV preferred to III) cannot be rationalised as the choices of a precise expected utility maximiser. The gambles are summarised in the table.

[2] We should note here that the evidence on peoples' actual choices in Ellsberg type set ups is multifarious and very far from univocal. Binmore, Stewart, and Voorhoeve (2012) find less ambiguity aversion than is often claimed. Voorhoeve et al. (2016) find that ambiguity aversion correlates with irresolution, a finding that suggests Moss' framework should account for this sort of behaviour. See Trautmann and Van der Kuijlen (2016) for an overview.

Table 1: The Ellsberg bets. The urn contains 30 red marbles and 60 blue/yellow marbles

	R	B	Y
I	1	0	0
II	0	1	0
III	1	0	1
IV	0	1	1

Let the probabilities for red, blue and yellow marbles be r,b and y respectively. If you were an expected utility maximiser and preferred I to II, then $r > b$ and a preference for IV over III entails that $r+y < b+y$. No numbers can jointly satisfy these two constraints. Therefore, no probability function is such that an expected utility maximiser with that probability would choose in the way described above.[3]

One might worry that the Ellsberg preferences might be due to *indifference* between the bets coupled with ambiguity aversion as some sort of tie-breaker. There is ample psychological evidence (Curley, Yates, and Abrams 1986; Curley and Yates 1989) that this tie-breaker view of ambiguity aversion cannot explain experimental subjects' patterns of preference.[4] The intuitive appeal of the Ellsberg preferences also seems to be insensitive to the exact number of marbles in the urn. You may still have the Ellsberg preferences if the urn contained 89 or 91 marbles (keeping fixed 30 red). Such insensitivity could not be accounted for by a "precise probability plus tie-break" model of ambiguity aversion.

A potential criterion of adequacy for an imprecise credence view is that it can rationalise making the Ellsberg choices permissible. (It would, I think, be too strong to require that an imprecise epistemology made the Ellsberg preferences obligatory.)

Let's interpret the Ellsberg problem as a sequential decision, and call it the *diachronic Ellsberg problem*. That is, I offer you the two choices (I versus II, then III versus IV)

[3]Actually eliciting such preferences in practice brings up some interesting theoretical problems that we don't have time to discuss, but see: Oechssler and Roomets (2014); Bade (2015)

[4]Other interpretations of the data may be possible, this is not the place to discuss this.

in quick succession. To be clear, I am going to draw two marbles from the urn (with replacement) and which of I or II wins will be decided by the first draw, whether III or IV win by the second draw. Note that this is different from Elga's decision problem, where both choices depended on the *same* draw. Both decisions have to be made before the you learn the outcome of either draw, so no learning takes place in between gambles. So you will be offered the two choices outlined below in quick succession. Note that the events have been subscripted to indicate which draw from the urn the gambles depend on.

- First choice:
- I which wins you 1 if R_1 nothing otherwise
- II which wins you 1 if B_1 and nothing otherwise
- Second choice:
- III which wins you 1 if $\neg B_2$ nothing otherwise
- IV which wins you 1 if $\neg R_2$ and nothing otherwise

We shall say that the pattern of preferences "I preferred to II and IV preferred to III" evinces *diachronic ambiguity aversion*.

Let's use the natural imprecise credence for the Ellsberg set-up which has $P(R_i) = 1/3$ and $P(B_i) = P(Y_i) = [0, 2/3]$ for $i = 1, 2$. Let's go through our choice rules again. The difference between naive and sophisticated gamma-maximin doesn't make a difference in this case. The minimum expectations for I,II,III,IV are $1/3, 0, 1/3, 2/3$ respectively. So both rules recommend I and IV. Indeed, they make the Ellsberg choices not just permissible, but *mandatory*.

Identification theory cannot rationalise the Ellsberg choices. If you identify with some p such that I is preferred to II, then that p also makes III preferred to IV. And if you identify with a p that prefers II to I, then it also prefers IV to III. On Moss' view, it is permissible, though presumably rare, for you to switch which element of P you identify with between the I vs. II choice and the III vs. IV choice. Switching is *precisely* what would be required for Moss to explain stable patterns of Ellsberg preferences. Indeed, for such patterns of preference to be stable would require systematic switching of which

$p \in P$ is identified with, which is something Moss rules out. So it is precisely what made Moss' view attractive with respect to the Hammond/Elga problem that precludes it from being a good rule in the Ellsberg problem!

The supervaluationist view has it that it is indeterminate whether it is permissible to choose I over II, and indeterminate whether it is permissible to choose II over I. Likewise for the choice between III and IV. But no probability is such that it sanctions preferring I over II and IV over III. So it is determinately impermissible to have the Ellsberg preferences. Rinard's view, therefore, cannot rationalise ambiguity aversion. And again, it's precisely what made the rule work in the Hammond/Elga problem that makes it fail here.

Forward looking – the rule advocated by Sud (2014) – cannot sanction the Ellsberg choices because they do not form an admissible sequence: for every probability function, there is some other sequence (either I then III or II then IV) that has higher expectation, thus the Ellsberg choices are not admissible.

I should say something about the special case of $p(B) = \frac{1}{3}$. What I say here applies equally to identification approaches, supervaluation approaches and forward looking approaches, though I'll talk in terms of identification. An agent who identifies with credence $p(B) = \frac{1}{3}$ will be indifferent between I and II, and between III and IV. Such an agent might well *choose* I over II and IV over III, without having a (strict) preference for those actions. Is this sufficient for the theory to claim to rationalise ambiguity aversion? I think not. First, allowing an indifference doesn't seem enough to rationalise a genuine preference for risky over ambiguous gambles. Second, such choice behaviour will be unstable in a way genuine ambiguity averse preference shouldn't be. Consider bet II$'$, which wins $1+\epsilon$ if a blue marble is drawn. Arguably, for small enough ϵ, an agent with genuinely ambiguity averse preferences would still prefer I to II$'$. That is, in a sequence of pairwise choices between:

- I vs II
- III vs IV
- I vs II$'$

it should be permissible to choose I, IV, I. However, identification theory says that (barring systematic switching of p) if you choose I and then IV, you will choose II′. This is so since the only $p \in P$ that sanctions choices for I and for IV is the "indifference-making" $p(B) = \frac{1}{3}$ and for this p, expectation of II′ is higher than I. So the identification theory can't accommodate the robustness of ambiguity averse preference. Nor can the theories of Rinard and Sud, for similar reasons.

So none of these rules can rationalise diachronic ambiguity aversion. And it is precisely those features of the rules that allow them to succeed in the Hammond/Elga problem that trip them up in the Ellsberg case.

5 Responding to the counterexample

One response open to those whose decision theories can't accommodate ambiguity aversion is to deny that accommodating ambiguity aversion should be a criterion of adequacy for imprecise decision theory. I have no knock-down argument against this response, but let me note two things. First, as I said at the beginning of the last section, Ellsberg problems are often used as a motivation for imprecise credences, so to deny that imprecise decision theory need to accommodate them is to undermine the case for imprecise credences. This seems somewhat self-defeating. Second, the motivations for claiming ambiguity aversion is not rational are often motivations for doing without imprecision altogether. So it seems like someone who wanted to ignore this constraint on imprecise decision will have to walk a fine line between dismissing ambiguity aversion on the one hand, and defending imprecision on the other. Conversely those who want to use imprecise decision theories' flaws as arguments against imprecision (as Elga did for his problem) will have a hard time arguing that decision theory ought to be ambiguity averse but not imprecise!

Given the failure of several decision theories to accommodate the Ellsberg decisions – a pattern of decisions that was taken to motivate IP – where does that leave imprecise credences? One response would be to attempt to find some criticism of the example. A possibility is to draw attention to the difference in intuitions between the case where the two choices of bets depend on the same draw of a marble as compared to the case where the choices depend on different draws from the urn. The Hammond/Elga problem is a case where it is important to the dialectic that both bets depend on the same outcome: whether or not H. The intuition that the Ellsberg choices are rational seems, to me, stronger when we consider the case where the two choices depend on different draws from the same urn. Perhaps there is room for manouevre here: those who want to salvage decision rules like those discussed above can downplay the importance of the example by emphasising the difference between sequences of choices depending on the same token of some unknown chance process (as in Hammond/Elga) versus sequences depending on distinct tokens of the same type of unknown chance (as in Ellsberg). I leave it to advocates of such decision rules to flesh out this line of response.

6 Where does this leave imprecise credence?

Recall that Adam Elga's original goal was to argue that no imprecise decision theory was adequate, and thus that rational subjective credence ought to be precise. Does the above demonstrate that Elga's project might still be brought to fruition? It seems that those rules that overcome the Hammond/Elga decision problem stumble when it comes to the diachronic versions of the Ellsberg problem. So again, perhaps no imprecise decision theory is adequate. I think such a conclusion would be premature.

First, note that the sophisticated maximin rule passes both tests. It, however, suffers from other problems. For example, it has issues with the value of information (Bradley and Steele 2016).^[

This is essentially because maximin does not satisfy the independence axiom (Seidenfeld 1988; Wakker 1988; Al-Najjar and Weinstein 2009).] A deeper issue with maximin is that it is too precise. One motivation for imprecise credences is to allow us to represent agents who don't have complete preferences, to allow for suspension of judgement (Seidenfeld 1993; Seidenfeld, Schervish, and Kadane 1995). So to then force imprecise decision theory to deliver a complete relation of preference among the options – as maximin does – seems a step backwards.

It seems like it would be a good thing if our imprecise decision theory allowed for incommensurability of options.[5] It is well known that incommensurability of goods leads to difficulties with diachronic consistency (Chang 1997; Broome 2000). The problem that Hammond and Elga were pointing towards is a problem of this type. So there is room for discussion about what precisely a response to the Hammond/Elga decision problem should involve. Bradley and Steele (2014) suggest that what should be ruled out is obligatory sure loss: that is, they argue that one should not have to rule out merely permissible sure loss. Hedden (2015) also seems to be arguing along these lines, in suggesting that "diachronic tragedy" does not entail irrationality.[6]

Second, note that no *precise* decision rule passes the diachronic Ellsberg test either, so if that test is plausible as a criterion of adequacy, then no precise decision theory is adequate either.

Here's an analogy:[7] the relationship between precise and imprecise credence is like that between propositional and first-order logic. There's a sense in which the former is a special case of the latter, but that undersells the novelty of the latter: there is an increase in expressive power in moving from the former to the latter. The increase in expressive power in using sets of probability functions comes in the form of, for example, distinctions between many independence concepts (Cozman 2012), and notions of symmetry (de Cooman and Miranda 2007) that collapse in the precise framework; the ability to have an adequate formal concept of ignorance (Norton 2007; 2008; 2010);[8] and many other things (see Bradley (2014) or Walley (1991), chapters 1 and 5 for further discussion). Complaining that decision theory has become harder and messier is like rejecting first order logic in favour of propositional logic because first order logic is

[5] Identity theory seems to determine complete preferences too, but Moss' explicit goal is to model *dilemmas* which seem to be characterised by the difficulty of making a choice, so presumably she would want to downplay or explain away the appearance of complete preferences in her theory.

[6] A diachronic tragedy is when you have beliefs and desires such that you perform a sequence of actions that you foresee will lead to a globally suboptimal outcome.

[7] This is one I owe to Greg Wheeler, in conversation

[8] Norton himself is sceptical of imprecise credence's prospects in this regard, but see Benétreau-Dupin (2015).

not decidable: yes, it's unfortunate, but it's something we have to deal with in order to make use of this superior – more expressive – formal framework. Imprecise credences and precise credences are no more in competition than are propositional and first order logic. Often, propositional logic is sufficient for our purposes, and we should use tools suitable to the problems we are interested in. Likewise, modelling agents as (precise) probabilists is often sufficient for our purposes. But sometimes there will be advantages to making use of a richer framework: examples include Benétreau-Dupin (2015), Singer (2014), Dorr (2010) (though Dorr is sceptical of the application)...

If you are certain about the true state of the world, you can act so as to maximise *actual* utility. If you only have knowledge of the objective chances, you can act to maximise *objective expected* utility. If you don't even have that, but have evidence that prompts you to have precise credences, you can act so as to maximise *subjective* expected utility. So we're comfortable with the idea that your epistemic state affects what is achievable in terms of decision making success. It's not surprising, then, that the goal-posts might shift again when we move to imprecise credence. What exactly should count as success in an imprecise decision context is, I think, still an open question.

7 Conclusion

Several attempts were made to overcome a problem for imprecise credences that was emphasised by Adam Elga. Three prominent such attempts (those of Moss (2015), Rinard (2015), and Sud (2014)) fail to get the right answers in another sequential decision problem, the diachronic Ellsberg problem. Although some might take this as an argument against imprecise credence, I see it instead as a spur to further work on the question "what can and what can't we say about rationality in circumstances of severe uncertainty?".

8 Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 792292.

Research also supported by the Alexander von Humboldt foundation and the University of Tilburg TiLPS research centre.

9 References

Al-Najjar, Nabil I., and Jonathan Weinstein. 2009. “The Ambiguity Aversion Literature: A Critical Assessment.” *Economics and Philosophy* 25: 249–84.

Bade, Sophie. 2015. “Randomization Devices and the Elicitation of Ambiguity-Averse Preferences.” *Journal of Economic Theory* 159: 221–35.

Benétreau-Dupin, Yann. 2015. “The Bayesian Who Knew Too Much.” *Synthese* 192: 1527–42.

Binmore, Ken, Lisa Stewart, and Alex Voorhoeve. 2012. “How Much Ambiguity Aversion? Finding Indifferences Between Ellsberg’s Risky and Ambiguous Bets.” *Journal of Risk and Uncertainty* 45: 215–38.

Bradley, Seamus. 2014. “Imprecise Probabilities.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. `http://plato.stanford.edu/archives/win2014/entries/imprecise-probabilities/`.

Bradley, Seamus, and Katie Steele. 2014. “Should Subjective Probabilities Be Sharp?” *Episteme* 11: 277–89.

———. 2016. “Can Free Evidence Be Bad? Value of Information for the Imprecise Probabilist.” *Philosophy of Science* 83: 1–28.

Broome, John. 2000. “Incommensurable Values.” In *Well-Being and Morality: Essays in Honour of James Griffin*, edited by R. Crisp and B. Hooker, 21–38. Clarendon Press.

Camerer, Colin, and Martin Weber. 1992. “Recent Developments in Modeling Preferences: Uncertainty and Ambiguity.” *Journal of Risk and Uncertainty* 5: 325–70.

Chandler, Jacob. 2014. “Subjective Probabilities Need Not Be Sharp.” *Erkenntnis* 79: 1273–86.

Chang, Ruth, ed. 1997. *Incommensurability, Incomparability and Practical Reason*.

Harvard University Press.

Cozman, Fabio. 2012. "Sets of Probability Distributions, Independence and Convexity." *Synthese* 186: 577–600.

Curley, Shawn P., and J. Frank Yates. 1989. "An Empirical Evaluation of Descriptive Models of Ambiguity Reactions in Choice Situations." *Journal of Mathematical Psychology* 33: 397–427.

Curley, Shawn P., J. Frank Yates, and Richard A. Abrams. 1986. "Psychological Sources of Ambiguity Avoidance." *Organizational Behavior and Human Decision Processes* 38: 230–56.

de Cooman, Gert, and Enrique Miranda. 2007. "Symmetry of Models Versus Models of Symmetry." In *Probability and Inference: Essays in Honor of Henry E. Kyburg Jnr.*, edited by William Harper and Gregory Wheeler, 67–149. Kings College Publications.

Dorr, Cian. 2010. "The Eternal Coin: A Puzzle About Self-Locating Conditional Credence." *Philosophical Perspectives* 24: 189–205.

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: 643–96.

Fox, Craig R., and Amos Tversky. 1995. "Ambiguity Aversion and Comparative Ignorance." *Quarterly Journal of Economics* 110: 585–603.

Gärdenfors, Peter, and Nils-Eric Sahlin. 1982. "Unreliable Probabilities, Risk Taking and Decision Making." *Synthese* 53: 361–86.

Halpern, Joseph Y. 2003. *Reasoning About Uncertainty*. MIT press.

Hammond, Peter. 1988. "Orderly Decision Theory." *Economics and Philosophy* 4: 292–97.

Hedden, Brian. 2015. "Options and Diachronic Tragedy." *Philosophy and Phenomenological Research* XC: 423–51.

Joyce, James M. 2010. "A Defense of Imprecise Credence in Inference and Decision."

Philosophical Perspectives 24: 281–323.

Levi, Isaac. 1974. "On Indeterminate Probabilities." *Journal of Philosophy* 71: 391–418.

———. 1980. *The Enterprise of Knowledge*. The MIT Press.

———. 1986. "The Paradoxes of Ellsberg and Allais." *Economics and Philosophy* 2: 23–53.

Moss, Sarah. 2015. "Credal Dilemmas." *Noûs* 49: 665–83.

Norton, John. 2007. "Probability Disassembled." *British Journal for the Philosophy of Science* 58: 141–71.

———. 2008. "Ignorance and Indifference." *Philosophy of Science* 75: 45–68.

———. 2010. "Cosmic Confusions: Not Supporting Versus Supporting Not." *Philosophy of Science* 77: 501–23.

Oechssler, Jörg, and Alex Roomets. 2014. "Unintended Hedging in Ambiguity Experiments." *Economics Letters* 122 (2): 243–46.

Rinard, Susanna. 2015. "A Decision Theory for Imprecise Probabilities." *Philosophers' Imprint* 15: 1–16.

Sahlin, Nils-Eric, and Paul Weirich. 2014. "Unsharp Sharpness." *Theoria* 80: 100–103.

Seidenfeld, Teddy. 1988. "Decision Theory Without 'Independence' or Without 'Ordering'. What's the Difference?" *Economics and Philosophy*, 267–90.

———. 1993. "Outline of a Theory of Partially Ordered Preferences." *Philosophical Topics* 21: 173–88.

Seidenfeld, Teddy, Mark J. Schervish, and Joseph B. Kadane. 1995. "A Representation of Partially Ordered Preferences." *Annals of Statistics* 23: 2168–2217.

Singer, Daniel J. 2014. "Sleeping Beauty Should Be Imprecise." *Synthese* 191: 3159–72.

Sud, Rohan. 2014. "A Forward Looking Decision Rule for Imprecise Credence." *Philosophical Studies* 167: 119–39.

Trautmann, S., and G. Van der Kuijlen. 2016. "Ambiguity Attitudes." In *Blackwell Handbook of Judgement and Decision-Making*, edited by G. Keren and G. Wu, Second, 89–116. Blackwell.

van Fraassen, Bas. 1990. "Figures in a Probability Landscape." In *Truth or Conse-*

quences, edited by Michael Dunn and Anil Gupta, 345–56. Kluwer.

Voorhoeve, Alex, Ken Binmore, Arnaldur Stefansson, and Lisa Stewart. 2016. "Ambiguity Attitudes, Framing and Consistency." *Theory and Decision* 81 (3): 313–37.

Wakker, Peter. 1988. "Nonexpected Utility as Aversion to Information." *Journal of Behavioural Decision Making* 1: 169–75.

Walley, Peter. 1991. *Statistical Reasoning with Imprecise Probabilities*. Vol. 42. Monographs on Statistics and Applied Probability. Chapman; Hall.