

A MODIFIED SUPERVALUATIONIST FRAMEWORK FOR DECISION-MAKING

Jonas KARGE

ABSTRACT. How strongly an agent believes in a proposition can be represented by her degree of belief in that proposition. According to the orthodox Bayesian picture, an agent's degree of belief is best represented by a single probability function. On an alternative account, an agent's beliefs are modeled based on a set of probability functions, called imprecise probabilities. Recently, however, imprecise probabilities have come under attack. Adam Elga claims that there is no adequate account of the way they can be manifested in decision-making. In response to Elga, more elaborate accounts of the imprecise framework have been developed. One of them is based on supervaluationism, originally, a semantic approach to vague predicates. Still, Seamus Bradley shows that some of those accounts that solve Elga's problem, have a more severe defect: they undermine a central motivation for introducing imprecise probabilities in the first place. In this paper, I modify the supervaluationist approach in such a way that it accounts for both Elga's and Bradley's challenges to the imprecise framework.

KEYWORDS: formal epistemology, supervaluationism, imprecise probabilities, decision-theory

Introduction

How strongly an agent believes in a proposition can be represented by her degree of belief in that proposition. According to the orthodox Bayesian picture, an agent's degree of belief is best represented by a single probability function. In particular, the Bayesian claims that agents must assign numerically precise probabilities to every proposition that they can entertain. On an alternative account, an agent's beliefs are modeled based on imprecise probabilities. With that, imprecise degrees of belief can be represented by a set of probability functions. A central decision-theoretical motivation for introducing imprecise probabilities is their solution to the Ellsberg Problem. In this problem, the orthodox Bayesian framework fails to adequately model the aversion of seemingly rational agents towards ambiguous actions whereas decision rules based on imprecise probabilities can do so.

Recently, however, imprecise probabilities have come under attack. Adam Elga¹ claims that there is no adequate account of the way they can be manifested in decision-making. In response to Elga, more elaborate accounts of the imprecise framework have been developed. One of them is based on supervaluationism, originally, a semantic approach to vague predicates. Supervaluationism can very naturally be applied to imprecise probabilities. With that, it solves Elga's problem.

Still, Seamus Bradley² showed that some of those accounts that solve Elga's problem, including supervaluationism, have a more severe defect: they undermine a central motivation for introducing imprecise probabilities in the first place. That is, their solution to the Ellsberg Problem. In this paper, I modify the supervaluationist approach in such a way that it accounts for both Elga's and Bradley's challenges to the imprecise framework.

This paper is organized as follows: In section 1, I will lay out the basic terminology of orthodox Bayesianism and the imprecise probabilities framework. In section 2, I will introduce the Ellsberg problem as a decision-theoretical motivation for imprecise probabilities. In section 3, we will look at Elga's decision-theoretical counterargument to imprecise probabilities. In section 4, I will show how supervaluationism can be applied to imprecise probabilities as well as how it solves Elga's problem. Moreover, I will introduce Bradley's argument against supervaluationism. In section 5, I will present a modified version of supervaluationism which solves Bradley's as well as Elga's problem.

1. Basic Terminology

In this section, I will briefly outline the basic terminology of the two competing views: Namely, *orthodox Bayesianism* and the *imprecise probabilities framework*.

1.1 Orthodox Bayesianism

The starting point for both views is an agent with beliefs about the world and who is capable of decision-making. According to the orthodox Bayesian picture, a *belief* can be defined as follows:

Definition, Belief. A belief is a ternary relation between an agent *S*, an object of belief and a real number between 0 and 1.³

¹ Adam Elga, "Subjective Probabilities should be Sharp," *Philosopher's Imprint* 10, 5 (2010): 1-11.

² Seamus Bradley, "A Counterexample to Three Imprecise Decision Theories," *Theoria* 85 (2019): 18-30.

³ Franz Huber, "Belief and Degrees of Belief," in *Degrees of Belief*, eds. Franz Huber and Christoph Schmidt-Petri (Dordrecht: Springer, 2009), 1-33, here 2.

We assume the objects of belief to be propositions, i.e. sets of possible worlds.⁴ Moreover, the real number assigned to a proposition by an agent is called her *degree of belief* in that proposition where 0 represents the lowest level of confidence and 1 the highest level of confidence in it. Consider, for instance, the proposition P that *it will rain tomorrow*. Assume, moreover, that agent S is 70% sure that it will, in fact, rain. We can then state that S's degree of belief in P is 0.7.

Bayesianism claims that degrees of belief ought to be represented by single probability functions that assign a precise number to propositions. In order to define probability functions, we have to be more precise about what we mean by *propositions*. For that purpose, we begin by defining the set of all possible worlds and calls this set *event space*:

Definition, Event Space. $\Omega = \{w_1, w_2, \dots, w_n\}$ is called an event space where each w_i in Ω is a state of affairs, or possible world.⁵

Since propositions are taken to be sets of possible words, we can define a proposition as follows:

Definition, Proposition. A proposition (or event) A is a subset of set Ω .⁶

Taking a proposition as a subset of the set of possible worlds, we can, moreover, define the set of all the propositions the agent can possibly believe in:

Definition, The Set of Propositions. The set of objects of beliefs (the propositions) is the power set of Ω : 2^Ω .⁷

Finally, based on our definition of the set of propositions, we can define a probability function as follows:

Definition, Probability Function. A probability function pr is a function $Pr: 2^\Omega \rightarrow \mathbb{R}$, satisfying the probability axioms.⁸

In a next step, we can apply this view to decision-making. When an agent has to choose between different actions, Bayesianism suggests as decision rule to choose the action that yields the highest expected utility. Thus, we not only need a precise

⁴ *Ibid.*, 2.

⁵ Anna Mahtani, "Imprecise Probabilities," in *The Open Handbook of Formal Epistemology*, ed. Richard Pettigrew and Jonathan Weisberg (PhilPapers Foundation, 2019), 107-130, here 108.

⁶ *Ibid.*, 108.

⁷ Seamus Bradley, "How to Choose Among Choice Functions," in *Proceedings of the Ninth Symposium on Imprecise Probability: Theories and Applications*, ed. Thomas Augustin, Serena Doria, Enrique Miranda, and Erik Quaeghebeur (2015), 57-66, here 57, <http://www.sipta.org/isipta15/data/paper/9.pdf>.

⁸ *Ibid.*, 57.

probability function but also a precise utility function that assigns a number to each possible outcome of an action reflecting the agent's value of that outcome.⁹ Given those functions, the expected utility for an action can be calculated as follows:

Let $\text{Pr}(S)$ be the degree of belief for an event to be the case, let $u(O)$ be the utility value the agent assigns to the consequence of an action, given the event S . Let A_i be some action. Now, the expected utility of A_i can be calculated as follows:

Definition, Expected Utility. $EU(A_i) = \sum_{j=1}^m \text{Pr}(S_j) \times u(O_{ij})$.

That is, we multiply the agent's degree of belief in an event by the utility value of the outcome of that action and, subsequently, sum those values for all possible outcomes.

To sum up, *orthodox Bayesianism* has two characteristics relevant to our discussion:

- 1) An agent's belief state is represented by a probability function. The probability function maps each relevant proposition to a real number between 0 and 1. This number is the agent's degree of belief in that proposition.
- 2) Rational agents must choose an action that has maximum expected utility based on the agent's degrees of belief in the relevant propositions.

1.2 Imprecise Probabilities

Assume, our agent S has to evaluate the proposition *the European Union will consist of exactly 27 member states in 20 years*. What precise probability should she assign to that proposition? Since orthodox Bayesianism represents an agent's belief with a single probability function, such a precise value has to be given.¹⁰ Considering propositions of this type, it seems highly implausible to represent belief states with a single probability function.¹¹

On an alternative account, degrees of belief can be defined based on imprecise probabilities. One way to construe imprecise probabilities is the following:

Definition, Imprecise Probabilities. Imprecise probabilities are sets of probability functions.¹²

⁹ Mahtani, "Imprecise," 119.

¹⁰ Susanna Rinard, "A Decision Theory for Imprecise Probabilities," *Philosopher's Imprint* 15, 7 (2015): 1-16, here 1.

¹¹ Miriam Schoenfield, "Chilling out on epistemic rationality," *Philosophical Studies* 158 (2012): 197-219, here 199.

¹² Seamus Bradely and Katie Steele, "Should Subjective Probabilities be Sharp?," *Episteme* 11, 3 (2014): 277-289, here 277.

Moreover, we call each such set of probability functions the agent's *representor* \mathcal{P} .¹³ Additionally, we assume that the set of values the distributions in the agent's representor assign to a proposition covers all of the interval $[x, y]$ with $x, y \in \mathbb{R}$. With that, we can define an imprecise degree of belief:

Definition, Imprecise Degree of Belief. An agent's imprecise degree of belief in a proposition H is represented by a representor, \mathcal{P} , with $\mathcal{P} = \{\Pr(H) : \Pr \in \mathcal{P}\}$.¹⁴

This can be illustrated as follows: Let A be the proposition that the European Union will consist of exactly 27 member states in 20 years. Assume, our agent S is 40-60% confident that this will be the case. With that, we can represent the agent's imprecise degree of belief in A with: $\mathcal{P}(A) = [0.4, 0.6]$.

Finally, an agent's representor can be understood as a credal committee where every probability function in that committee represents the opinion of one of its members.¹⁵ Collectively, these opinions reflect the beliefs of an agent.¹⁶ This idea will be significant for the last section of this paper.

2. The Ellsberg Problem

A central decision-theoretical motivation for introducing imprecise degrees of belief is the so-called Ellsberg Problem. For this problem, it is vital to distinguish between risky and ambiguous actions. We take an action to be risky in case the probabilities of the relevant outcomes are known. An action is ambiguous, in turn, if the probabilities are unknown, or, only partially known.¹⁷ The Ellsberg Problem relies on the observation that, under specific circumstances, seemingly rational agents prefer taking risky decisions, instead of ambiguous ones, even though it violates expected utility theory.¹⁸

¹³ *Ibid.*, 227.

¹⁴ Bradley, "How to Choose," 57.

¹⁵ Seamus Bradley and Katie Steele, "Learning, and the 'Problem' of Dialation," *Erkenntnis* 79 (2014): 1287-1303, here 1291.

¹⁶ Seamus Bradley, "Imprecise Probabilities," *The Stanford Encyclopedia of Philosophy* (Spring 2019 Edition), <https://plato.stanford.edu/archives/spr2019/entries/imprecise-probabilities>.

¹⁷ Bradley, "A Counterexample," 22.

¹⁸ Katie Steele, "Distinguishing Indeterminate Belief from 'Risk-Averse' Preferences," *Synthese* 158 (2007): 189-205, here 190.

2.1 The Ellsberg Problem with Precise Probabilities

Now, to the problem itself: In the Ellsberg Problem, an agent is told that an urn contains 30 red balls and 60 balls that are either blue or yellow in some unspecified proportion.¹⁹ The agent faces two decision problems: A and B.

In problem A, the agent can decide to bet on either (I) which yields \$100 if the next ball drawn is red or (II) where she receives \$100 if it is blue. Likewise, in problem B: The agent can bet on (III) where she gets \$100 if the next ball drawn is not blue or (IV) she receives \$100 if it is not red.²⁰ For the sake of simplicity, we can assume that receiving \$100 yields a utility value of 1 and not receiving it yields 0 utility.²¹ With that, we can summarize the payoffs as follows:²²

	Red	Blue	Yellow
Problem A			
(I)	1	0	0
(II)	0	1	0
Problem B			
(III)	1	0	1
(IV)	0	1	1

Figure 1. Payoffs, Ellsberg.

A significant majority of apparently rational people chooses option (I) in problem A and option (IV) in problem B when the Ellsberg Problem is studied empirically.²³ In the following, we will call this combination of (I) and (IV) *Ellsberg preferences*. Since the probabilities for those actions are known to the agent, they classify as risky actions. By upholding to this pattern, the agent expresses an aversion towards the ambiguous options (II) in problem A and (III) in problem B.

The orthodox Bayesian, however, cannot rationalize the Ellsberg preferences since there is no precise probability an agent could possibly assign to drawing a blue ball such that $EU(I) > EU(II)$ and, at the same time, $EU(III) < EU(IV)$.²⁴ This can be

¹⁹ *Ibid.*, 191.

²⁰ *Ibid.*, 191.

²¹ Daniel Ellsberg, "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics* 75, 4 (1961): 643-669, here 655.

²² Figure based on: Bradley, "A Counterexample," 23.

²³ Steele, "Distinguishing," 191.

²⁴ *Ibid.*, 191.

shown as follows: Let $EU(I) = Pr_1$, $EU(II) = Pr_2$, $EU(III) = Pr_1 + Pr_2$, and $EU(IV) = Pr_2 + Pr_3$. However, there is no Pr_i such that $Pr_1 > Pr_2$ and $Pr_1 + Pr_3 < Pr_2 + Pr_3$.²⁵

As ambiguity aversion seems to be a feature of rational decision-making, not being capable of adequately modeling it is a problem for the Bayesian account.

2.1 The Ellsberg Problem with Imprecise Probabilities

Now, let's analyze the Ellsberg Problem with imprecise probabilities. Since the proportion of blue and yellow balls is unknown, we can assume the proportion to lie somewhere between 0 and 2/3. It could be the case, for instance, that all non-red balls turn out to be blue or that there are as many blue balls as yellow balls. Thus, a natural distribution of probability functions is the following:

Imprecise Degrees of Belief, Ellsberg. $\mathcal{P}(\text{blue}) = \mathcal{P}(\text{yellow}) = [0, 2/3]$ and $Pr(\text{red}) = 1/3$.²⁶

Representing an agent's belief state with imprecise degrees of belief implies that the expected utility for a given action will be imprecise also. That is, it corresponds to a set of utility values given by the probability functions in the agent's representor. Since those can overlap, the possible actions under considerations can turn out to be incommensurable.²⁷ Given the imprecise probabilities for the Ellsberg problem, the expected utilities can be summarized as:

Imprecise Expected Utilities, Ellsberg. (I) = 1/3, (II) = [0, 2/3], (III) = 1/3 + [0, 2/3], (IV) = 2/3.²⁸

Since expected utility theory cannot be applied to intervals, we have to look for alternative decision rules for imprecise probabilities.²⁹ One such possible rule is the *Maximin Rule*. It tells us the following: For each action, there is a lowest expected utility value. This lowest value is the minimum expected utility for an action. In a decision problem, Maximin recommends the agent to choose the action that has the maximum minimum expected utility.³⁰

In order to apply Maximin to the Ellsberg Problem, we first have to state its minimum expected utility values:

Minimum expected utility values, Ellsberg. $EU(I) = 1/3$, $EU(II) = 0$, $EU(III) = 1/3$, and $EU(IV)$

²⁵ Ellsberg, "Risk, Ambiguity," 655.

²⁶ Steele, "Distinguishing," 195.

²⁷ Bradley and Steele, "Should Subjective," 278.

²⁸ Steele, "Distinguishing," 195.

²⁹ Nils-Eric Sahlin, "Unsharp," *Theoria* 80, 1 (2014): 100-103, here 100.

³⁰ Mahtani, "Imprecise," 121.

Jonas Karge

$= 2/3$.

With that, the Maximin Rule recommends choosing (I) in problem A and (IV) in problem B. This corresponds to the Ellsberg preferences. Hence, we have shown that there is at least one decision rule for imprecise probabilities that rationalizes the Ellsberg preferences, and, with that, it rationalizes a case of ambiguity aversion.

To sum up, empirical studies indicate that agents tend to have the Ellsberg preferences even though it violates expected utility theory. According to orthodox Bayesianism, those agents are irrational. However, it seems like they are, in fact, rational.³¹ With that, we have a case where orthodox Bayesianism fails to adequately model an instance of rational decision-making: it does not succeed in modeling ambiguity aversion. Still, the Ellsberg preferences can be rationalized with the imprecise framework.³² This is taken to be a decision-theoretical motivation for the imprecise probabilities framework.

3. Elga's Problem

Even though decision rules based on imprecise probabilities seem to perform well in cases of ambiguity aversion, they struggle in another type of decision problem. That is, when imprecise probabilities are applied to sequential decision problems.³³ When it comes to sequential decision problems, it can be the case that each decision taken individually is rationally admissible, the sequence of decision, however, can turn out to be rationally impermissible.³⁴

The central sequential decision problem to this discussion has been presented by Adam Elga. Elga's argument is structured as follows: He considers three types of possible decision rules for imprecise probabilities and shows for each type that it either leads to absurd consequences in a specific decision problem, or, that it has some other severe defect. For this discussion, we will only consider the first type of decision rule.

The specific decision problem goes as follows: In a *great series of bets*, an agent is sequentially offered Bet A first, and, immediately after the agent decides whether to accept or reject Bet A, she is offered Bet B.³⁵ Now, let H be some proposition such as *it will rain tomorrow*. The agent is then offered the following series of bets:

Bet A. If H is true, S loses \$10. Otherwise S wins \$15.

³¹ Steele, "Distinguishing," 190.

³² Mahtani, "Imprecise," 125.

³³ *Ibid.*, 191.

³⁴ Bradley, "A Counterexample," 21.

³⁵ Elga, "Subjective," 4.

Bet B. If H is true, S wins \$15. Otherwise S loses \$10.

This series of bets is called *great* since S is guaranteed to win \$5 in case she accepts both bets.³⁶ However, it is not rationally required for S to accept both bets. If she believes, for instance, that H is highly unlikely, she would be better off to accept Bet B only. Still, she is rationally required to accept at least one of the bets since rejecting both bets is strictly dominated by accepting them.³⁷ With that, any decision rule that permits to reject both bets rationalizes an obviously irrational action, and, thus, has to be rejected.

Let's apply Maximin to the great series of bets. The minimum expected utility for rejecting Bet A is 0. However, the one for accepting Bet A is -10. With that, Maximin suggests rejecting Bet A. Likewise, the minimum expected utility for rejecting B is 0 whereas the one for accepting it is -10.³⁸ Hence, our agent S should reject both bets according to Maximin. This, however, is irrational and Maximin does fail in this betting scenario.

Finally, this result can be extended to a number of decision rules for imprecise probabilities which all suggest rejecting both bets by Isaac Levi, Peter Walley, Teddy Seidenfeld, Gärdenfors and Sahlin as well as Gilboa.³⁹

4. Supervaluationism

A defender of imprecise probabilities in decision-making can now choose one of two strategies:

Strategy 1. It can be argued that the great series of bets does not show that decision rules for imprecise probabilities are irrational.

Strategy 2. It can be accepted that it does, but that different decision rules for imprecise probabilities can be introduced that are not affected by Elga's argument.⁴⁰

In this section, I will now introduce an approach following strategy 2: namely, supervaluationism. Supervaluationism is, originally, a semantic theory designed to handle vague predicates. The central idea is that vague predicates such as *tall* don't have a definite extension, but rather a variety of different extensions. Each possible extension of a vague predicate corresponds to a possible precisification of that

³⁶ *Ibid.*, 4.

³⁷ *Ibid.*, 4.

³⁸ Bradley, "A Counterexample," 20.

³⁹ Elga, "Subjective," 5.

⁴⁰ Richard Pettigrew, *Dutch Book Arguments (Draft)* (Cambridge University Press: 2019), 96, <https://richardpettigrew.com/books/the-dutch-bookargument/>.

predicate.⁴¹ Since there is no definite precisification for the semantics of a vague language, the semantic value of a statement remains unclear unless there is complete agreement among the precisifications on that value.⁴²

When it comes to the truth value of statements in a vague language, *complete agreement* is understood as a proposition being either *determinately true* or *determinately false* in supervaluationistic terms. This can be spelled out as follows:⁴³

Definition, Determinately True. If a proposition is true according to all admissible precisifications, then it is determinately true.

Definition, Determinately False. If a proposition is false according to all admissible precisifications, then it is determinately false.

That said, if neither of those two is the case, it is possible for a statement to have no semantic value:

Definition, Indeterminately True. If a proposition is true according to some, but not all, admissible precisifications, then it is indeterminate whether it's true.

Example. Consider the predicate *tall*. For this predicate, there are numerous possible precisifications. Each such precisification determines a threshold for what it means to be tall and not-tall. Assume, every threshold between 160cm and 200cm is an admissible precisification, but 220cm is not an admissible precisification. In that case it's determinately true that someone 220cm tall is tall whereas it's indeterminately true whether someone 170cm tall is tall.⁴⁴

In a next step, we have to make more precise what we mean by an *admissible* precisification. In fact, supervaluationism can very naturally be applied to imprecise probabilities by giving such a precisification:

Definition, Admissible Precisifications. The admissible precisifications are the functions in an agent's representor.⁴⁵

Moreover, with that definition at hand, we can characterize a supervaluationist decision theory based on imprecise probabilities. Contrary to most decision theories, actions will now not only be permissible or impermissible, but also classifiable as indeterminately permissible. This can be seen from the following definitions:⁴⁶

⁴¹ Rosanna Keefe, "Vagueness: Supervaluationism," *Philosophy Compass* 3, 2 (2008): 315-324, here 315.

⁴² Achille C. Varzi, "Supervaluationism and Its Logics," *Mind* 116 (2007): 633-676, here 634.

⁴³ Definitions according to: Rinard, "A Decision Theory," 2.

⁴⁴ *Ibid.*, 2.

⁴⁵ *Ibid.*, 2.

⁴⁶ Definitions according to: Rinard, "A Decision Theory," 3.

Definition, Determinately Permissible Action. If some action A has the highest expected value (or ties for highest) according to every function in the agent's representor, then it's determinately true that A is permissible.

Definition, Determinately Impermissible Action. If some action A has a higher expected value according to every function in the agent's representor than some alternative action B, action B is determinately impermissible.

Analogous to the evaluation of the semantic value of statements in the supervaluationist framework, we also have the case of indeterminate permissibility:

Definition, Indeterminately Permissible. If some action A has the highest expected value according to some, but not all, functions in the representor, it is indeterminate whether A is permissible.

4.1 Supervaluationism and Sequential Decision-Making

In a next step, we can apply these definitions to Elga's problem. Assume, again, the great series of bets:

Bet A. If H is true, S loses \$10. Otherwise S wins \$15.

Bet B. If H is true, S wins \$15. Otherwise S loses \$10.

Assume, moreover, the following representor for our agent S as suggested by Elga: $\mathcal{P}(H) = [0.1, 0.8]$.⁴⁷

Now, according to our supervaluationist decision theory, it is indeterminate whether accepting Bet A is rationally permissible. This is the case because it is permissible according to some, but not all probability functions in the agent's representor. For instance, according to the function that represents the precise value of $\Pr(H) = 0.2$, accepting Bet A has the highest expected value. According to $\Pr(H) = 0.7$, though, the agent should reject Bet A. Likewise, it is indeterminate whether rejecting Bet B is permissible. However, it is determinately impermissible to reject both bets since there is an alternative action with a higher expected value according to every function in the agent's representor.⁴⁸ This alternative is to accept both bets. Since this analysis yields the desired result, the supervaluationist decision theory does succeed in Elga's problem.

4.2 The Sequential Ellsberg Problem

Even though supervaluationism succeeds in Elga's problem, Bradley recently showed that it fails in another: If we interpret the Ellsberg Problem sequentially,

⁴⁷ Elga, "Subjective," 4.

⁴⁸ Rinard, "A Decision Theory," 6.

supervaluationism cannot rationalize ambiguity aversion. If this is the case, supervaluationism undermines a central motivation for introducing imprecise probabilities in the first place.⁴⁹ For this reason, Bradley's argument is a major challenge for the supervaluationist account.

Let's see how supervaluationism fails in the sequential Ellsberg Problem. As in the original problem, we again have an urn that contains 90 balls. From those balls, 30 balls are red and the remaining ones are either blue or yellow in some unknown proportion. Now, according to the sequential interpretation, an agent is offered two decision problems in quick succession:

Problem A. The agent faces two choices:

(I), which wins the agent a utility value of 1 if the ball drawn in the first round is red and nothing otherwise.

(II), which wins the agent a utility value of 1 if the ball drawn in the first round is blue and nothing otherwise.

Problem B. The agent faces two choices:

(III), which wins the agent a utility value of 1 if the ball drawn in the second round is not blue and nothing otherwise.

(IV), which wins the agent a utility value of 1 if the ball drawn in the second round is not red and nothing otherwise.⁵⁰

As degree of belief for our agent, we assume that $\mathcal{P}(\text{blue}_i) = \mathcal{P}(\text{yellow}_i) = [0, 2/3]$ and $\text{Pr}(\text{red}_i) = 1/3$ with $i = 1, 2$ referring to the current round of the decision problem.⁵¹ With that, we can begin by analyzing Problem A.

Analysis, Problem A. According to supervaluationism, it is indeterminate whether it is permissible to choose (I) over (II) and (II) over (I). It is indeterminate to choose (I) over (II) because it is permissible according to those functions in the representor that assign a probability less than $1/3$ to drawing a blue ball and impermissible to the other functions in the representor. Likewise, it is permissible according to those functions in the representor to choose (II) over (I) that assign a probability greater than $1/3$ to drawing a blue ball, but impermissible according to the other functions. The same line of reasoning applies to problem B:

Analysis, Problem B. It is indeterminate whether it is permissible to choose (III) over (IV) and (IV) over (III).

⁴⁹ Bradley, „A Counterexample,“ 18.

⁵⁰ *Ibid.*, 24.

⁵¹ *Ibid.*, 24.

So far, every option is indeterminately permissible. But what about the Ellsberg preferences, i.e. the sequence of (I) in round 1 and (IV) in round 2?

Analysis, Ellsberg Preferences (I) + (IV). No function in the representor is such that it yields a preference of (I) over (II) and (IV) over (III). With that, it's determinately impermissible to have the Ellsberg preferences.⁵²

This is a major drawback for a supervaluationist decision theory for imprecise probabilities. Initially, the Ellsberg Problem was used in order to motivate imprecise probabilities since there are decision rules for them that can rationalize ambiguity aversion. Still, we have now shown that supervaluationism fails to do so.⁵³ With that, supervaluationism seems to undermine a central motivation for introducing imprecise probabilities in the first place.

Concluding this section, I want to briefly discuss the idea of *rationalizing* the Ellsberg preferences. To begin with, it is clear that any decision rule that classifies them as irrational, or determinately impermissible in this case, fails.

However, it remains unclear how much permissibility is necessary in order to rationalize the Ellsberg preferences. One option is to classify them as determinately permissible. However, this could be a too strong requirement. That is, we are looking for the right amount of permissibility that we should be assigning to them.⁵⁴ In fact, it can be argued that classifying them as indeterminately permissible is just the right amount. To support this idea, I want to give the following line of reasoning:

On the one hand, we do not want to rationalize the Ellsberg preferences by classifying them as determinately permissible since there is no precise probability that allows this pattern of preferences. Thus, in order to respect expected utility theory, this amount of permissibility is too much.

On the other hand, as we have seen, we want to take seriously the aversion towards ambiguous gambles among rational agents. Therefore, it would be too strong to classify the Ellsberg preferences as determinately impermissible.

Luckily, supervaluationism allows for a third class of actions. Thus, in order to solve this conflict, I introduce a modified version of supervaluationism that classifies the Ellsberg preferences as indeterminately permissible as what I take to be the right amount of permissibility.

⁵² *Ibid.*, 25.

⁵³ *Ibid.*, 25.

⁵⁴ The idea of the right amount permissibility regarding the Ellsberg preferences goes back to personal correspondence with Dr. Seamus Bradley.

5. Modified Supervaluationism

In this section, I introduce a modified supervaluationist framework which has to meet two objectives: First, it has to rationalize the Ellsberg preferences in its diachronic version. Secondly, it must not rationalize the rejection of both bets in Elga's problem.

The starting point for the modified supervaluationist framework is to take literally the interpretation of the agent's representor as a credal committee. In this committee, I construe every member as a voter that votes for propositions or actions in decision problems. As voting method, I apply relative majority voting. That is, the alternative that accumulates the most votes wins.⁵⁵ Based on this idea, I will now modify the central concepts of supervaluationism as well as its decision rule.

According to the original supervaluationist account, a proposition is *determinately true* if it is true according to all admissible precisifications. In the following, I replace this notion by propositions being *predominantly true*. This can be defined as follows:

Definition, Predominantly True. A proposition is predominantly true if it is true according to a relative majority of precisifications.

Consider the following example: Assume, that there are ten admissible precisifications for the predicate *tall*. According to one of those the threshold for being tall is 170cm, according to two precisifications it is 175cm, according to three it is 185, and according to four precisifications the threshold lies at 180cm. In this case, it is predominantly true that someone who is at least 180cm tall is tall.

In a next step, we apply this idea to imprecise probabilities by assuming that every probability function in the representor is represented by a member of a voting committee. With that, we can very naturally derive a novel decision rule for imprecise probabilities. We begin by defining *predominantly permissible* and *impermissible* actions:

Definition, Predominantly Permissible Action. An action A is predominantly permissible if a relative majority of members in the representor vote for it.

Definition, Predominantly Impermissible Action. An action A is predominantly impermissible iff there is a relative majority of members in the representor that vote for an alternative action B.

⁵⁵ Joachim Behnke, Florian Grotz, and Christof Hartmann, *Wahlen und Wahlsysteme* (Oldenbourg: De Gruyter, 2016), 8.

It is important to note that an available action does only count as impermissible in case there is an alternative that receives a majority of votes. If this is not the case, we have a case of indeterminacy:

Definition, Determinately Permissible Action. If there is no relative majority in the representor for any action, every action is indeterminately permissible.

Finally, it has to be defined how the members of the committees do, in fact, vote:

Definition, Vote for an Action. A member in the representor, representing a probability function, votes for the action with the highest expected utility. If all actions bear the same expected utility according to the probability function, this member refrains from voting.

Consider the following example:

Example, Vote for an Action. Let action U be: buy an umbrella; action (I): buy ice cream. Furthermore, let H be: It will rain in an hour with $\mathcal{P}(H) = [0.1, 0.6]$.

	H	$\neg H$
U	1	-1
I	-1	1

Figure 2. Payoffs, Vote for an Action.

The members in the representor which represent $\mathcal{P}(H) = [0.1, 0.5]$ vote for (I) (80%). The members which represent $\mathcal{P}(H) = (0.5, 0.6]$ vote for U (20%). With that, (I) is predominantly permissible and U is predominantly impermissible.

5.1 Modified Supervaluationism and the Sequential Ellsberg Problem

In the final part of this text, I will first apply the modified supervaluationist framework to the sequential Ellsberg Problem, and, subsequently, to Elga's problem.

The agent is facing again problem A and B where she has to choose between (I) and (II) as well as (III) and (IV) respectively in two rounds. Moreover, we have as imprecise degrees of belief: $\mathcal{P}(\text{blue}) = \mathcal{P}(\text{yellow}) = [0, 2/3]$ and $\text{Pr}(\text{red}) = 1/3$. With that, we can analyze both problems as follows:

Analysis, Problem A. The members in the voting committee that represent $\mathcal{P}(\text{blue}) = [0, 1/3]$ vote for (I) since it yields a higher expected utility than voting for (II).

The ones representing $\mathcal{P}(\text{blue}) = (1/3, 2/3]$, in turn, vote for option (II). The member that represents the function with $\text{Pr}(\text{blue}) = 1/3$ refrains from voting. With that, there is no majority for either option, and, thus, both options are indeterminately permissible. The same holds for problem B:

Analysis, Problem B. $\mathcal{P}(\text{blue}) = [0, 1/3)$ vote for (III) and $\mathcal{P}(\text{blue}) = (1/3, 2/3]$ vote for (IV). Both of these actions are indeterminately permissible.

Finally, since this is the sequential Ellsberg Problem, we have to consider the sequence of (I) in round 1 and (IV) in round 2. That is, the Ellsberg preferences. Similar to the original supervaluationist framework, no member in the representor votes for this sequence. However, according to modified supervaluationism, that does not imply its impermissibility.

Analysis, Ellsberg preferences. The members in the representor that represent the probability functions $\mathcal{P}(\text{blue}) = (1/3, 2/3]$ vote for (II) + (IV) and the ones that represent $\mathcal{P}(\text{blue}) = [0, 1/3)$ vote for (I) + (III). With that, the Ellsberg preferences are indeterminately permissible.

This, I count as an advantage of modified supervaluationism because it confirms with the demanded *right amount of permissibility* that should be assigned to the Ellsberg preferences by neither classifying them as determinately permissible nor determinately impermissible.

5.2 Supervaluationism and Elga's Problem

In a final step, I will apply modified supervaluationism to Elga's problem. Our agent faces once more the great the series of bets:

Bet A. If H is true, S loses \$10. Otherwise S wins \$15.

Bet B. If H is true, S wins \$15. Otherwise S loses \$10.

We assume, moreover, the imprecise degree of belief in H given by Elga: $\mathcal{P}(H) = [0.1, 0.8]$. Now, we can analyze both bets as follows:

Analysis, Bet A and B. The members in the representor representing $[0.1, 0.6)$ vote to accept Bet A. The members representing $(0.6, 0.8]$ vote to refuse Bet A. With that, 71% vote to accept Bet A.

Likewise, Bet B: The members representing $[0.1, 0.4)$ vote to refuse Bet B and the ones representing $(0.4, 0.8]$ to accept Bet B. With that, 57% vote to accept Bet B.

Thus, accepting Bet A and accepting Bet B are predominantly permissible. For the given representor this is the correct result. Moreover, this can be shown for any possible representor:

Assertion: For no imprecise degree of belief, it is possible to refuse both bets.

Proof. Assume, that it is possible to refuse both bets. In particular, the agent has then to refuse Bet A. In order to refuse Bet A, the agent has to have more functions in his representor with $\Pr(H) > 60\%$. This is the case for imprecise degrees of belief with $\mathcal{P} \subseteq (0.6, 1]$. Bet B, in turn, is voted to be accepted for any imprecise degree of belief with $\Pr(H) > 40\%$. That is: $\mathcal{P} \subseteq (0.4, 1]$. With that, every member that represents a function with $\Pr(H) > 60\%$ votes for A to be refused but for Bet B to be accepted. Thus, it is not possible to refuse both bets at the same time.

Summary

This paper's objective was to provide a decision-theoretical framework based on imprecise probabilities that solves Elga's and Bradley's challenge. By modifying the supervaluationist account such a framework could be found. Modified supervaluationism construes the agent's representor as a voting committee that applies relative majority voting to evaluate the truth of statements and permissibility of actions. Moreover, it relies on a weaker notion of truth and permissibility than standard supervaluationism. Instead of *determinate* truth and permissibility, modified supervaluationism only requires *predominant* truth and permissibility. With that, it succeeds in both cases: It rationalizes the Ellsberg preferences to a reasonable extend and it does not rationalize rejecting both bets in Elga's problem.