1 Introduction

A college graduate deciding whether to enter medical school may not assign precise probabilities or precise utilities to the possible outcomes of her options. She may not evaluate precisely her chances of succeeding in medical school or the attraction of life as a physician. A person making a choice in a decision problem often does not have enough information or experience to assign precise probabilities and utilities to her options' possible outcomes. Common principles of rational choice, such as the principle to maximize expected utility, cover only decision problems with precise probabilities and utilities for possible outcomes. However, a general account of rational choice must also cover decision problems with imprecise probabilities and utilities for possible outcomes.

An explanation of imprecise probabilities and utilities includes as a foundation a thorough account of precise probabilities and utilities. I take probabilities as rational degrees of belief and utilities as rational degrees of desire and in support of this interpretation argue that rational degrees of belief comply with the axioms of probability and that rational degrees of desire comply with the principle of expected utility. Then I extend this view to imprecise probabilities and utilities. Afterwards, I advance a decision principle that uses imprecise probabilities and utilities to identify rational choices.

The decision principle applies within a model of choice that makes several idealizations about agents and their decision problems. The principle lays the groundwork for more general principles that dispense with the idealizations.

My treatment of imprecision briefly discusses rival positions but does not thoroughly review the literature. References direct readers to alternative stances.

Section 2 characterizes imprecise probabilities and utilities and Section 3 defends their rationality. Sections 4 and 5 present constraints that rationality imposes on precise probabilities and utilities and Section 6 extends the constraints to imprecise probabilities and utilities. Section 7 formulates a principle of rational choice that accommodates imprecise probabilities and utilities. Section 8 applies the principle in sequences of choices and Section 9 applies the principle to choices in games of strategy. Section 10, the final section, draws conclusions.

2 Imprecision

An account of imprecise probabilities and utilities grows out of an account of precise probabilities and utilities. This section explains what probabilities and utilities are and then what it means for probabilities and utilities to be imprecise. Later sections explain how to use imprecise probabilities and utilities to make decisions. The principle of rational choice I present relies on this section's points about agents, their decision problems, and their resources for resolving their decision problems.

2.1 A Decision Model

I advance a principle of rational choice in a decision model – that is, under a set of assumptions about agents and their decision problems. The principle uses an agent's belief states, or doxastic states, and the agent's desire states, or conative states. The model idealizes agents but assumes that they have doxastic and conative states of the type that humans have. Psychology describes the types of mental states that humans have and philosophical points about mental states help define and individuate them. I use a lay understanding of psychology to describe doxastic and conative states, attending especially to features philosophically important for the model's principle of rational choice.

The decision model incorporates several idealizations and simplifying assumptions about agents, their circumstances, and their decision problems. In the model, agents are cognitively ideal and so know all a priori truths they entertain – but need not entertain every a priori truth even though reflection is effortless and instantaneous for them. Furthermore, ideal agents know their own minds and so are aware of their doxastic and conative states and their cognitive powers. I assume that, in the decision problems they face, they are aware of the relevant characteristics of the problems and resolve the problems without distraction. They frame their decision problems using sentences that express their options and the possible outcomes of their options. I assume that they have information sufficient to understand fully the proposition any sentence expresses (in the language they use, with at least the expressive power of English). Therefore, they recognize when two sentences express the same option or the same possible outcome of an option.

Lacking evidence concerning a proposition's truth, an ideal agent need not assign a precise probability to the proposition, even given unlimited reflection. An agent's reflection on the proposition and the grounds of a probability assignment to the proposition cannot make up for a lack of information. An imprecise probability may remain imprecise, despite reflection, if new information does not arrive. Similarly, a utility that is imprecise because of a lack of experience may remain imprecise, despite reflection, if new experience does not arrive. Reflection cannot make up for a lack of experience.

A probability is imprecise if no single number accurately characterizes it. In this case, a number characterizing it is indeterminate. The same holds for an imprecise utility. Its imprecision entails the indeterminacy of a number characterizing the utility.

Because a proposition's probability is a number representing an agent's doxastic attitude to the proposition, its indeterminacy characterizes a representation of the attitude and not the attitude itself. The attitude has exactly the features it has. The same point applies to the indeterminacy of a proposition's utility. A proposition's utility is a number representing an agent's conative attitude to the proposition and its indeterminacy characterizes a representation of the agent's conative attitude. The attitude, given exactly its features, has an indeterminate numerical representation. Moreover, even if a doxastic or conative attitude has an indeterminate representation by a single number, an alternative type of representation may accurately characterize the attitude.

I say that an attitude is *quantitative* if it has an apt representation that uses a single quantity, such as a number. If an attitude is nonquantitative, an apt representation may still use numbers, but not a single number, to represent the attitude. A representation of the attitude that uses an interval of numbers is quantitative although it does not specify a single number to represent the attitude. It imprecisely specifies a number to represent the attitude.

When evaluating a proposed principle of rationality, I assume that an agent in the decision model is rational in all respects except possibly for compliance with the principle and then consider whether rationality requires compliance with the principle. An argument for the principle rescinds the idealization that the agent is rational in matters the principle governs. Then it maintains that an agent's violation of the principle is irrational.

I take a *doxastic domain* for an agent to be a set of propositions to which the agent has doxastic attitudes.[1] These attitudes may yield probability assignments but may be just judgments of epistemic possibility or impossibility. The set is usually taken to be a *Boolean algebra* formed from a set of atomic propositions by closure under negation and disjunction. However, to allow for multiple equivalent propositions differing in structure, I take the set to form a *propositional language* constructed from a set of atomic propositions by closure under the standard propositional operations.[2]

[1] An agent's adopting a doxastic domain for probability assignments prevents inconsistencies that may arise if she assigns probabilities to all propositions.

[2] An agent may fail to assign probabilities to propositions in the doxastic domain she uses not because of imprecision but because she regards some propositions as infinitesimally less probable than others. I put aside such cases by assuming that the agent uses a doxastic domain that is *Archimedean* in the sense that, for any two epistemically possible propositions compared probabilistically, there is a natural number n such that one proposition is no more than n times more probable than the other.

In an agent's decision problem, the number of *salient* options is finite just in case in some adequate representation of the decision problem the options form a finite set such that the agent is not indifferent between any two options belonging to the set. In this case, a unique option maximizes utility; ties do not arise. A representation of the options in a decision problem may combine tying options, options such that the agent is indifferent among them, into a single option; the option may be the disjunction of the tying options. Realizing a disjunctive option is equivalent to realizing a disjunct.

Using this terminology, the decision problems that the decision model treats have the following characteristics. The number of salient options is finite and options have a finite number of possible outcomes with probabilities and finite, stable utilities not altered by an option's adoption. If the probabilities and utilities of possible outcomes are imprecise, they nonetheless have an adequate representation.

An agent's probability assignments may use various doxastic domains. In the decision model, suppose that an agent adopts a doxastic domain for a decision problem, assigns probabilities to propositions in the domain that express possible outcomes of options, and reaches a choice. The agent's adopting another doxastic domain for the decision problem does not reverse the choice. Each doxastic domain yields the same choice. So, although a doxastic domain's selection is arbitrary, its arbitrariness is inconsequential. An agent's doxastic domains differ in the events to which they assign probabilities but do not differ, for propositions to which they assign probabilities, in features that affect choice.

If a rational ideal agent ideally situated in a decision problem that the decision model treats assigns precise probabilities and utilities to the possible outcomes of options, an option is rational if and only if it maximizes expected utility. Section 7 generalizes this principle to cover decision problems in the model without precise probabilities and utilities. A further extension of the principle beyond the model may follow common methods of generalizing a principle in a model by relaxing idealizations and revising the principle to accommodate new situations.

2.2 Probability

Probabilities come in two sorts. *Physical probabilities* arise from physical features of events, such as the physical features of a coin toss. *Evidential probabilities* are relative to evidence and an agent's evidential probabilities are relative to the agent's evidence. Imagine a courtroom trial of a defendant for commission of a crime. The physical probability that the defendant is guilty depends on the defendant's past acts and is either 0 or 1, depending on whether the defendant committed the crime. For a juror, the evidential probability that the accused is guilty depends on the evidence presented during the trial and may rise or fall as evidence is presented even though the physical probability of guilt is constant. At the end of the trial, the evidential probability of guilt may have a non-extreme value although the physical probability of guilt remains either 0 or 1. Physical probabilities are sometimes called objective but probabilities may be objective in many ways. For example, an evidential probability may be objectively settled by the evidence. Evidential probabilities are sometimes called epistemic but they are epistemic in a specific way; they are relative to evidence.

A rational ideal agent knows the evidential probabilities of an option's possible outcomes relative to her evidence; and her degrees of belief, which direct her choices, equal these evidential probabilities. Principles of rational choice using probabilities use evidential probabilities because an agent often does not know the physical probabilities of an option's possible outcomes. Because my topic is rational choice, by probability I generally mean evidential probability.

An evidential probability attaches to a proposition, the content of a declarative sentence. I take a proposition to have a structure similar to the structure of a sentence. Two logical truths may therefore have different structures and so be different propositions, although each is true in all possible worlds. Because two distinct propositions may be true in exactly the same possible worlds, a set of possible worlds does not adequately represent a proposition. Sentences and the propositions they express are similar, so I sometimes speak of the probability of a sentence, meaning the probability of the proposition that the sentence expresses.

I understand events in a technical sense that includes states. Propositions represent events such as acts, states of the world that settle the consequences of acts, and the outcomes of acts. Probabilities and utilities attach to events by attaching to propositional representations of the events.[3]

The term imprecise probability is a bit misleading because a probability in the ordinary sense is precise. Indeterminate probability is a more suggestive term.[4] However, I use the term imprecise probability, taking it in a technical sense that does not entail being a probability, because this usage is widespread. An imprecise probability is an imprecise specification of a probability, such as an interval of probabilities.[5]

3 Jeffrey ([1965] 1990) attaches probability and utilities (desirabilities) to propositions but takes propositions to be adequately represented by sets of possible worlds.
4 Levi (1974) recommends using the term indeterminate probability.
5 Walley (1991) provides a classic account of imprecise probabilities. Bradley (2019) and Mahanti (2019) offer recent surveys. Augustin et al. (2014), Troffaes and de Cooman (2014), and Zaffalon and Miranda (2018) present mathematical results.

A person may assign probabilities to some, but not all, propositions of a doxastic domain she adopts. Her probability that heads turns up on a coin toss may be precisely 50 percent. However, because of sparse information, she may not assign a precise probability to rain tomorrow. Imprecision concerning the atoms of a doxastic domain may spread to compounds formed from the atoms. For example, imprecise probabilities for two atoms may lead to an imprecise probability for their disjunction.

Familiar doxastic states include belief, suspension of judgment, and disbelief. A *degree of belief* quantitatively represents a doxastic state. The relation between belief and degree of belief clarifies degree of belief. However, the relation is subtle. A belief is not simply a high degree of belief, as the threshold for what counts as a high degree of belief must at least vary with context to accommodate beliefs about lottery tickets – an agent typically has a very high degree of belief that any given lottery ticket will lose but still does not believe that the ticket will lose and instead suspends judgment. The doxastic attitudes that degrees of belief represent explain beliefs but in a complex way. Assuming that an agent is both cognitively ideal and rational simplifies an account of the relation between belief and degree of belief because the assumption puts aside cases in which, for example, an agent irrationally believes a proposition to which she assigns a low degree of belief. However, even for a rational ideal agent, the relation between belief and degree of belief is intricate. I do not describe the relation except to say that belief that a proposition is true is generally the product of a high degree of belief that the proposition is true, according to a context-sensitive threshold for being high. I do not need to be more specific about the relation because I treat in detail only degree of belief, and not belief, and so do not need a detailed, unified account of doxastic attitudes.

I take a doxastic attitude that a degree of belief represents as a primarily passive response to evidence (such as observation) and not as a primarily active representation of the world (so that it is evaluable as an act). The doxastic attitude is sometimes called a strength of belief, with the understanding that minimum strength is not belief but disbelief. An agent's degree of belief that one proposition holds is greater than the agent's degree of belief that another proposition holds only if the agent believes the first proposition more strongly than the second proposition, again with the understanding that, when the two degrees of belief are low, the agent typically does not believe either proposition but instead disbelieves both. Degree of belief has a technical sense according to which it does not measure only belief but also disbelief. Some authors use the term credence instead of degree of belief to disavow a restriction to belief.

Degrees of belief represent doxastic attitudes. A degree of belief that a proposition holds is a number representing an agent's doxastic attitude to the proposition. The number represents strength of belief (in a technical sense that includes strength of disbelief), with, by convention, 1 representing maximum strength and 0 representing minimum strength. Because degrees of belief are numbers, they can comply with the laws of probability. I take evidential probabilities as rational degrees of belief.

Representations of attitudes differ from the attitudes themselves. A degree of belief is a number and not itself a doxastic attitude. For brevity of expression, theorists sometimes speak of a degree of belief as if it were the attitude it represents. They say, for example, that a person's high degree of belief that a proposition is true explains the person's belief that it is true. Strictly speaking, they mean that the doxastic attitude, the strength of belief, that the high degree of belief represents explains the person's belief. For convenience, I also sometimes speak of a degree of belief as if it were a doxastic attitude, although strictly speaking it is a number representing a doxastic attitude.

By convention, degrees of belief, as I understand them, represent ratios of strengths of belief not just differences in strengths of belief. Hence, degrees of belief use a ratio scale rather than, say, an interval scale, so that, if an agent's degree of belief that *p* is 0.6 and the agent's degree of belief that *q* is 0.3, then the agent believes *p* twice as strongly as *q*.

An account of degree of belief may implicitly define it by advancing principles governing it that are sufficient for grasping the meaning of degree of belief. Probabilism, the view that rational degrees of belief comply with the laws of probability, a view that Section 4 advocates, contributes to an account of degree of belief that implicitly defines it using, among other principles, the laws of probability as norms for degrees of belief. Psychological descriptions of the causes of strengths of belief, such as evidence, and the effects of strengths of belief, such as acts, further supplement this section's brief introduction of degrees of belief.

2.3 Utility

Conative attitudes include desire, indifference, and aversion. The conative attitudes that degrees of desire represent are quantitative counterparts of desire, indifference, and aversion. Although an agent may form a conative attitude, such as a desire, putting aside some considerations, degrees of desire represent attitudes formed all things considered.

As is degree of belief, degree of desire is a technical term. Using indifference as a zero point for a scale of degree of desire, a negative degree of desire represents an aversion and a positive degree of desire represents a desire.[6]

[6] In one sense, indifference holds between two items; an agent is indifferent between them. In another sense, indifference is toward a single item. An agent may be indifferent to dessert.

Accordingly, the relation between desire and degree of desire is less complex than the relation between belief and degree of belief.

A desire and a conative attitude represented by a degree of desire are distinct attitudes but may have the same realization. The same mental state may be classified as a quantitative attitude that a degree of desire represents and also classified as a nonquantitative desire, aversion, or attitude of indifference.

Desire and degree of desire and, similarly, aversion and degree of aversion (or negative degree of desire) are passive attitudes responding to events entertained (and are not evaluable as acts), although they prompt acts to satisfy desires and to prevent realizations of aversions. As for degree of belief, degree of desire has an implicit definition given by principles governing it, including the normative principle of expected utility that requires an act's degree of desire to equal the expected degree of desire of the act's outcome. An account of the causes of strengths of desire, such as envy, and the effects of strengths of desire, such as acts, fills out the implicit definition.

The conative attitudes that degrees of desire represent are propositional attitudes, that is, attitudes directed toward propositions, as are desire, indifference, and aversion. Because degrees of belief and degrees of desire alike attach to propositions, principles such as the expected-utility principle may join them seamlessly.

Degrees of desire that satisfy the laws of utility are called utilities, just as degrees of belief that satisfy the laws of probability are called probabilities. Utilities as well as probabilities may be imprecise. A person may assign a precise utility to gaining a thousand dollars but not assign a precise utility to holding a lottery ticket that if drawn yields a thousand dollars because she does not know the number of tickets in the lottery. Also, a person who lacks the experience of eating passion fruit may not assign a precise utility to eating this fruit. A conative attitude, such as a desire, may have an indeterminate numerical representation because of a lack of information or a lack of experience. For consistency of terminology, just as I use the term imprecise probability instead of indeterminate probability, I use the term imprecise utility instead of indeterminate utility. An imprecise utility is an imprecise specification of a utility, such as an interval of utilities.

Theorists often claim that an agent, even with reflection, may not have a preference, or be indifferent, between two events, such as hiking along Hurricane Ridge in the Olympic National Park and listening to a performance of Beethoven's Ninth Symphony by the Chicago Symphony Orchestra. She may

Indifference as a zero point is indifference in the second, non-relational sense. The two senses are closely related. An agent is indifferent to dessert if and only if she is indifferent between having dessert and not having it.

not be able to compare the events along any convenient dimension of comparison, not even pleasure, because the pleasures the events generate are of different types. If a rational ideal agent assigned both the hiking and the listening a precise degree of desire, then she would prefer the event with the higher degree of desire, and if their degrees of desire were the same would be indifferent between them. Because degrees of desire entail comparability, incomparability entails the absence of degrees of desire and so imprecision.[7]

For an ideal agent, degrees of desire, when rational, comply with the laws of utility, such as the expected-utility principle. Therefore, I take utilities as rational degrees of desire. When a rational ideal agent in a decision problem fails to assign degrees of desire to her options, the options have imprecise utilities for her.

2.4 Constructivism

My account of degrees of belief takes them to represent doxastic attitudes to propositions. It counts as a *realist* view, as opposed to a *constructivist* view that takes degrees of belief to represent not attitudes but choices or preferences and so to be constructed from choices or preferences rather than to have an independent reality. One constructivist account, following de Finetti ([1937] 1964), defines an agent's degree of belief that a proposition holds as the smallest percent of a dollar that the agent will exchange for a bet that gains a dollar if the proposition holds and otherwise nothing. The norm of expected-utility maximization, assuming that dollar amounts equal utilities, requires that the agent pay for the bet a percent of a dollar no greater than the agent's degree of belief that the proposition holds. This constructivist account makes the relation between the degree of belief and the exchange rate hold by definition, whereas my realist account accommodates the relation's being a normative requirement.

Another constructivist account defines an agent's degree of belief that a proposition holds as the value of the probability function that represents the agent's doxastic comparisons of propositions. This account assumes that the comparisons satisfy certain conditions, presented by Krantz and colleagues (1971: chap. 5), that ensure the existence and uniqueness of a probability function representing the comparisons. The definition makes having degrees of belief that satisfy the axioms of probability dependent on an arbitrary selection of a representation of doxastic comparisons, given that some perfectly adequate representations do not use a probability function, as, for example, Titelbaum (forthcoming: sec. 14.1) observes. In contrast, my realist account

7 Chang (1997) offers a collection of essays on incomparability.

takes satisfying the axioms of probability as a normative requirement for degrees of belief. It also strengthens the norms for doxastic comparisons of propositions: not only must the comparisons be representable as agreeing with a probability function but they must also agree with the particular probability function that rational degrees of belief form.

A third constructivist account, for both degrees of belief and degrees of desire, defines an agent's degree of belief that a proposition holds as the value of the probability function that, along with a utility function, represents the agent's preferences among gambles as following expected utilities. This account assumes that the preferences satisfy certain conditions, for example those presented by Savage ([1954] 1972), that ensure the existence and uniqueness of the probability function and the existence and uniqueness (up to a positive linear transformation) of the utility function that together represent the preferences. The definition makes having preferences among gambles that follow expected utilities dependent on an arbitrary selection of a representation of the preferences, because some perfectly adequate representations of the preferences do not have them follow expected utilities, as, for example, Titelbaum (forthcoming: sec. 8.3) observes.[8] In contrast, my realist account takes following expected utilities as a normative requirement for preferences among gambles. The requirement is not just that preferences among gambles be representable as following expected utilities; they must follow expected utilities, calculated using degrees of belief and degrees of desire that are defined independently of preferences among gambles, as in Subsections 2.2 and 2.3. These degrees of belief and degrees of desire uniquely represent an ideal agent's doxastic and conative state, assuming it is quantitative, given a scale for degrees of desire.

Defining probabilities and utilities using choices, so that choices maximize expected utility by definition, destroys the power of the principle of expected-utility maximization to explain the rationality of choices. If choices maximize expected utility by definition, their maximizing expected utility cannot explain their rationality, not even their meeting requirements for having a representation as maximizing expected utility. Although the principle to choose as if maximizing expected utility does not use expected utility to explain the rationality of choices, the stronger, traditional decision principle of expected-utility maximization does, assuming an interpretation of probabilities and utilities according to which they represent propositional attitudes. It may take probability as rational degree of belief and utility as rational degree of desire, given that degrees of belief and degrees of desire represent propositional attitudes and exist independently of choices. The traditional principle of expected-utility

8 Lyle Zynda (2000) and Meacham and Weisberg (2011) make similar observations.