Uncertainty in the Context of Pragmatist Philosophy and Rational Choice Theory

Jeffrey Helzner
Columbia Univeristy

December 11, 2011

1

What is meant by *uncertainty*? Presumably it is a term that has been used in a variety of ways. The present article focuses on some of the interesting relations between uncertainty and both pragmatist philosophy and rational choice theory. Section 2 is concerned primarily with the ways that uncertainty has been accommodated in the study of rational choice and, in particular, the ways in which attempts to accommodate uncertainty have motivated departures from the orthodox Bayesian tradition that finds its roots in the work of Ramsey and Savage. Section 3 is concerned primarily with the ways that the aforementioned departures from the orthodox Bayesian tradition can be seen as arising from a commitment to pragmatist principles of the sort that are developed in the work of Peirce and Levi.

2

For the purposes of the present article, the meaning of the term *decision theory* will be restricted to a cluster of topics that concern synchronic standards of single-agent rationality as reflected in principles of rational choice. The development of *expected utility theory* forms the backbone of the intended cluster of topics. The following three remarks are meant to further clarify the suggested restrictions:

First, we are restricting to single-agent standards, as opposed to standards of rationality concerning groups of agents that might interact with each other in various ways. This is not meant to suggest that there are no interesting developments relating to standards that are concerned with groups of agents, quite the contrary. Rather, it is because there is such great variety in these developments. Much of this variety results from the fact that so many different kinds of assumptions have been made concerning the relevant group of agents, e.g., that they are in competition with each other or perhaps that certain information is salient to every member of the group. As a result, many of these developments now constitute large and well-defined subfields of rational choice theory, e.g., *game theory* and theories of *social choice*. Note, however, that much of what will be considered in the present article is presupposed in discussions of multi-agent contexts. Thus, the reader is advised not to conclude that there is no significant interaction between decision theory and a topic like game theory [14, 34].

Second, we are restricting to synchronic standards of rationality. This is because developments concerning diachronic standards typically depend crucially on theories of updating – e.g., *Bayesian updating* – that can be studied somewhat independently of decision theory in connection with topics like the confirmation of hypotheses. Again, the reader should not conclude that the topics covered in the present article are unrelated to diachronic standards of rational choice, since these standards typically presuppose many of the synchronic principles that will be considered in what follows.

Third, we will be focusing on a particular expected utility tradition, one that is perhaps most closely associated with the work of Savage as presented in [36]. It seems clear that the tradition exemplified by Savage's work is the expected utility tradition that is most well-known to economists, psychologists, and statisticians. However, the situation in philosophy is a bit different, since philosophers often draw upon the expected utility tradition that is perhaps most clearly associated with the work of Jeffrey as presented in [12]. The present focus on the tradition associated with Savage is not meant to suggest that the tradition associated with Jeffrey is less significant. Rather, the choice of focus is in recognition of the fact that discussions of Jeffrey's approach are often held in the context of other well-developed topics in the philosophy of decision theory – e.g., debates between *evidential* and *causal* decision theories [13] – as well as philosophical topics that can be pursued somewhat independently of decision theory, e.g., the assignment of probability to conditionals. Once again, the reader is urged not to conclude that the present article is irrelevant to the philosophical tradition that has just been described, since many of the relevant discussions – e.g., [13] – are informed by the Savage tradition despite their departures from it.

Even with its scope restricted in the manner that has just been described, the study of decision making is relevant to a variety of fields – e.g., economics [7, 45], philosophy [12, 24], psychology [16, 6], and statistics [36, 15] – and it seems reasonable to assume that many people are interested in certain kinds of decisions – e.g., financial decisions – but these considerations do not imply that there are any interesting *theories* of decision. For it may be that all of the interesting aspects of decision making are so highly contextual that they resist non-trivial theorizing. Presumably, non-trivial theories must have some sort of relevance. There are at least three ways in which a given theory of decision might be relevant [18]. It might be *normatively* relevant, providing a standard with which decisions ought to agree. It might be *prescriptively* relevant, providing techniques to improve decision making. It might be *descriptively* relevant, providing a basis for explaining and predicting the behavior of decision makers. There are various potential dependencies among these three sorts of relevance:

It might be expected of a prescriptively relevant theory that it will help to bring decision makers in line with the demands of a certain normatively relevant theory. It might be expected of a normatively relevant theory that it be descriptively relevant in at least some situations, e.g., those involving sophisticated decision makers in situations where the "cognitive load" does not impose too much of a burden on the decision maker's computational capacities.

In addition to questions concerning the way in which a theory of decision might be relevant, there are questions concerning the structure of such a theory. For example, what is required in order for a thing to qualify as a presentation of such a theory? Perhaps there are various things - e.g., metaphors, diagrams, or equations - that could figure into such a presentation. Our focus in the present article is on "formal" theories. Such theories serve to determine a class of structures, e.g., the models of the theory. In this sense such theories may be viewed as part of ordinary mathematics, like group theory or topology. However, the theories at issue are not primarily of interest as pure mathematics. The central interest, as is typical in emerging field of *formal epistemology* [10, 11], concerns the relevance of these structures to the world in which decisions are made [42, 44]. A given formal theory might be relevant in more than one of the ways mentioned above, but some of these ways might be regarded as more compelling than others; e.g., optimization of transitive preferences might count as having some sort of normative relevance, even if, in light of empirical evidence, it is not regarded as having great descriptive relevance.

Of the three notions of relevance that have been mentioned, it seems clear that normative concerns have dominated philosophical discussions of decision theory. This is not surprising, since norms of decision making serve to circumscribe some important aspect of rationality, a traditional topic of interest for philosophers. Such interest in rationality is often predicated on the thesis that agents *ought* to be rational, where the term *rational* denotes some suitably articulated standard for individual agents; for example, taking the term to mean expected utility maximizer yields the idea that agents ought to be expected utility maximizers. But why make such an assumption? Although a detailed examination of this sort of question is not appropriate for the present survey, a few brief remarks are in order. Hence, one possible justification for such an assumption is that there are reasons to believe that it is true. Such a justification requires that statements like "agents ought to be rational" are truth-value bearing, a matter that is hardly clear: Is the truth-value of such a statement contingent upon the facts of the world? If so, then what can be said about the empirical content of a statement like "agents ought to be rational"? In contrast, if such statements are not contingent, then are they to be construed as necessary in some sense? If so, then what can be said about the principles of logic that will underwrite such necessity? This line of attack threatens to drag us into a Russian winter of old and difficult philosophical questions.

In contrast to that which was just considered, a rather different conception of the thesis at issue couches the discussion in terms of agent-relative notions rather than categorical ones. For example, it might be suggested that an agent who is committed to being rational is an agent who ought to be rational, on pain of failing to live up to its commitments. The idea that one should desire to live up to ones commitments is taken here as constitutive of the very concept of having a commitment. The suggested move to commitments, and reactions to it, might be compared to those that are associated with certain aspects of the pragmaticist tradition in epistemology, where, following Peirce, the "fixation of belief" is central and concerns with the "pedigree of knowledge" play a less substantial role [31, 24, 28]. In both contexts, an agent-relative concept – e.g., commitments or beliefs – is taken to provide a sufficient foundation for further investigations into rational choice in general and epistemology in particular.

To be sure, a move to commitments raises its own questions, with the first of these concerning the conditions that must be satisfied for there to be such commitments. One sort of answer to such questions maintains that whether or not a commitment is in place is a matter of psychological fact, although perhaps one that is only accessible through introspection. Savage's remarks concerning the status of his own subjective expected utility theory as presented in *The Foundations of Statistics*, which is perhaps still the most influential book on decision theory, are worth recalling in connection with answers of the sort that was just described:

> I am about to build up a highly idealized theory of the behavior of a "rational" person with respect to decisions. In doing so I will, of course, have to ask you to agree with me that such and such maxims of behavior are "rational." In so far as "rational" means logical, there is no live question; and, if I ask your leave there at all, it is only as a matter of form. But our person is going to have to make up his mind in situations in which criteria beyond the ordinary ones of logic will be necessary. So, when certain maxims are presented for your consideration, you must ask yourself whether you try to behave in accordance with them, or, to put it differently, how you would react if you noticed yourself violating them. [36]

The essence of Savage's position, as presented in the quoted passage, appears to be that the agent can introspect and determine if it is committed to the standard at issue. For example, following the suggestion made at the end of Savage's remarks, the agent has a commitment to a given standard at time t if at t the agent judges that it would seek appropriate therapy upon noticing itself violating that standard. Note that on such a view an agent who is committed to a given standard at t might not satisfy that standard at t – the committed agent is of course obliged to seek therapy when it notices that it has failed to live up to its commitment, but that is a different matter. For example, an agent that is committed to Savage's theory could very well have weak preferences that are incompatible with that theory. However, so long as that commitment is in place, such an agent is obliged to take steps towards eliminating those incompatibilities when they are recognized by that agent [24].

As discussed, decision theory, at least in its more formal branches, has been concerned with the problem of articulating standards of rational choice and, minimally, such standards serve to determine a class of structures. Let us now focus matters by considering the case of *optimization*, a standard that is assumed in much of contemporary decision theory. Roughly, an optimizing agent is one that, if offered a menu of alternatives, would constrain its selection to an alternative that it weakly prefers to all of the other alternatives on the given menu. Assuming a set of alternatives, and of which the menus of alternatives are subsets, and where weak preference on a menu is representable as a binary relation on that set, optimization might be seen as determining a subclass of those structures of type $\langle A, \mathcal{M}, \mathcal{P}, \mathcal{C} \rangle$ where A is a nonempty set, $\mathcal{M}$ is a subset of the set of all subsets of A, $\mathcal{P}$ is a function that associates each $M \in \mathcal{M}$ with a binary relation on M, and $\mathcal{C}$ is a function that associates each $M \in \mathcal{M}$ with a subset of M. The intended interpretation of these components being as follows: A is the set of *alternatives*; $\mathcal{M}$ is the set of all *potential menus*; and $\mathcal{P}$ assigns each menu to the binary relation that represents the agent's *weak preferences* on that menu. The final component, $\mathcal{C}$, assigns each menu to the set of *admissible alternatives* on that menu – we will return to this notion later, but the basic idea is that the admissible alternatives are those that can be selected in a way that satisfies the standard to which the agent is committed.

With the relevant type of structure determined, optimization can be formulated more precisely as follows: $\langle A, \mathcal{M}, \mathcal{P}, \mathcal{C} \rangle$ is an optimization structure just in case, for all menus M, $\mathcal{C}(M)$ is the set of all those alternatives x in M such that $(x, y) \in \mathcal{P}(M)$ for all y in M. This formulation of optimization is rather weak, and there is a tradition of restricting to cases in which additional requirements are met. For example, it is often required that $\mathcal{C}(M)$ is nonempty whenever M is nonempty – i.e., that each nonempty menu contain at least one admissible alternative – and this, of course, further constrains $\mathcal{P}$ in light of the quasi-formal statement of optimization that has just been given. In fact, there are at least two other general constraints that are imposed on $\mathcal{P}$ in typical formulations of optimization [38, 22, 35]. First, it is assumed that $\mathcal{P}(M)$ must be a weak order if it is to represent a rational agent's weak preferences over M [32]. Second, it is assumed that for rational agents $\mathcal{P}$ is given by restriction in the sense that if $M \subseteq M'$, then $\mathcal{P}(M)$ is the restriction of $\mathcal{P}(M')$ to M. This assumption of menu-independence is often strengthened so as to guarantee the existence of a weak order P on A such that, for each menu M, $\mathcal{P}(M)$ is the restriction of P to M [40].

It is well-known that under standard assumptions – roughly the two sorts of constraints just mentioned – the $\mathcal{P}$ component of an optimization structure is definable from its remaining components. For example, suppose that $\langle A, \mathcal{M}, \mathcal{P}, \mathcal{C} \rangle$ is an optimization structure that satisfies the two sorts of constraints that have been mentioned and for which $\mathcal{M}$ is the set of all subsets of a finite, nonempty A. It can be shown that for such an optimization structure the global preference P, which is assumed to represent $\mathcal{P}$ by virtue of satisfying the second constraint from the previous paragraph, can be recovered from $\mathcal{C}$ as $(x, y) \in P$ iff $x \in \mathcal{C}(\{x, y\})$. Hence, optimization structures can also be regarded as a subclass of those structures of type $\langle A, \mathcal{M}, \mathcal{C} \rangle$, where, as before, A is a nonempty set of alternatives, $\mathcal{M}$ is the set of all potential menus from A, and $\mathcal{C}$ is a function that associates each menu M with a set of admissible alternatives from

M. Under standard assumptions – e.g., $\mathcal{M}$ consists of all finite subsets of A – the optimization structures can be viewed as those structures of type $\langle A, \mathcal{M}, \mathcal{C}\rangle$ that satisfy Sen's α and β conditions. α requires that if M and M' are menus such that $M \subseteq M'$ and x is an alternative in M that is also admissible in M' – i.e., $x \in M \cap \mathcal{C}(M')$ – then x is admissible in M. β requires that if M and M' are menus such that $M \subseteq M'$ and both x and y are admissible alternatives in M, then either both of those alternatives are admissible in M' or neither of them are admissible in M' [38, 22].

Optimization is often taken as a necessary principle of rational choice, but it is seldom taken as a sufficient one. In particular, it seems reasonable to many that a rational agent is, minimally, one that is disposed to restrict its selection to those alternatives that, given its beliefs, best serve its desires, but it is clear that the more general conception of optimization does not impose such a requirement since the weak preferences being optimized are not required to have any connection with the agent's beliefs or desires. The requirement at issue needs to be made more precise if it is to do any heavy lifting. Various questions arise concerning what is meant by *beliefs* and *desires* as well as the sense in which some alternatives "best serve" relative to such beliefs and desires. The most influential answers to these questions have come from the *expected utility* tradition [22]. According to the simplest theories from this tradition an alternative may be thought of as a kind of lottery, not unlike those that are run by various states as a way funding social programs. Such an assumption about alternatives requires that if f is an alternative for an agent, then that agent knows the probability $p_f(i)$ of each possible outcome i that is associated with selecting f. If we are content with this assumption about alternatives and are willing to further restrict our attention to those agents whose judgments concerning the desirability of the various possible outcomes can be represented in terms of a real-valued utility function u, then the expected utility tradition offers a familiar standard of rationality by restricting to the class of optimization structures for which the underlying binary relation, P, is such that $(f, g) \in P$ just in case the expected utility of f is at least as great as that of g – that is, P, which is to be interpreted as a relation of weak preference, is such that

$$(f, g) \in P \text{ iff } \sum_{i \in I} p_f(i)u(i) \geq \sum_{i \in I} p_g(i)u(i) \tag{2.1}$$

for every pair of alternatives f and g.

It is worth pausing to consider some of what is presupposed in (2.1). First, as noted above, alternatives are lotteries. In particular, for each lottery f there is a probability distribution p_f such that, for each potential outcome $i \in I$, $p_f(i)$ is the probability that f yields i. Such a probability distribution is assumed to represent some objective feature of the underlying lottery structure through which it is determined. Those who are interested in the conceptual foundations of "objective" probability – e.g., *frequencies* versus *propensities* – are urged to consult [42]. A second presupposition in (2.1) is that both sums are well-defined. There are various technical options as far as how to satisfy this requirement. For example, one could restrict attention to lotteries that assign positive probability to only a finite number of prizes, in which case the sums at issue involve nothing more than adding together a finite number of terms. Alternatively, one could follow the familiar strategy of generalizing (2.1) through a theory of integration [22]. A third presupposition in (2.1) concerns the function u that is supposed to represent the extent to which the agent finds the possible outcomes desirable. Since the domain of u is required to be the set of possible outcomes, (2.1) presupposes that the desirability of a possible outcome does not depend on the lottery under which it might be awarded. A fourth presupposition is that this u representing the desirability of outcomes is required to be real-valued and appropriate relative to the taking of sums and products. In particular, the scale-type of this representation must be restricted enough to underwrite the *meaningfulness* of those arithmetic operations on u that are involved in (2.1) [41, 21, 29]. For example, if u, as a representation of the agent's judgments of desirability, is merely unique up to ordinal-transformations, then the expected utility calculations in (2.1) fail to be meaningful in the sense that they depend on an unmotivated choice of representation. Hence, the utilities in (2.1) are assumed to be *cardinal* utilities, with a particular representation like u being unique up to a positive linear transformation – in this sense it is presupposed in (2.1) that utility, like temperature, is measureable on an interval scale.

Arguments concerning the normative relevance of the expected utility hypothesis embodied in (2.1) have been offered in terms of the "law of large numbers" and "self-evident" axioms [32]. Apart from reactions to examples such as those offered by Allais [1], there seems to be wide consensus for expected utility maximization as the appropriate standard of rationality for decision making under *risk*, i.e., those instances of decision making where the agent knows, for each alternative, the probabilities associated with the various possible outcomes of selecting that alternative. A significant limitation of this standard is that it is often the case that the decision maker does not know all of the relevant probabilities.

Cases in which the agent does not know all of the relevant probabilities are said to involve *uncertainty* rather than mere risk. While various decision rules have been suggested for decision making under uncertainty - e.g., security-oriented rules like *maximin* that received significant attention in connection with cases of "complete ignorance" [26] – the most influential of these extend the expected utility tradition to uncertainty by introducing subjective, or personal, probabilities in the absence of known, objective probabilities. The introduction of subjective probabilities raises several interesting questions in the philosophy of probability. For example, what are these subjective probabilities supposed to represent? The standard answer, following [33] and [2], is that they represent *credal* states, or degrees of belief. Assuming that rational agents have credal states, why should every potential state of that sort be representable as a probability function? In response to such a question it is customary to cite the *Dutch book argument* [33] that purports to show that agents who have credal states that violate the suggested standard can be exploited by a clever bookie. It is worth noting that the significance of such arguments are not only limited by their assumptions [17] but also by the fact that they can be adapted to support rival positions [43]. Finally, what sorts of propositions are allowed to be in the domain of a subjective probability? Can a rational agent assign a degree of belief – as supposedly underwritten by the Dutch book argument – to its own potential acts [12]? Can a rational agent assign such degrees of belief to subjunctive conditionals [13]? Further discussion of these questions, though interesting they might be, is beyond the scope of the present survey.

Perhaps the most influential arguments against the assumption at issue – i.e., that every potential credal state of a rational agent can be represented by a numerically precise probability distribution – locate the inadequacy of that assumption in connection with the role that credal states are supposed to play within extensions of expected utility theory to cases of uncertainty. To appreciate these arguments one must consider how alternatives are understood in the usual extensions of expected utility theory to decision making under uncertainty. According such extensions, such as in the account offered by Savage, an alternative is a function from the set of possible states of nature to the set of possible consequences [36]. If we are willing to restrict our attention to those agents who have a subjective probability $\mathcal{P}$ on the events of some distinguished partition, I, of the set of possible states of nature, and a cardinal utility u over the set of possible consequences, then the previously considered expected utility hypothesis is readily extended to the class of class of optimization structures for which the underlying binary relation, P, is such that $(f, g) \in P$ just in case the subjective expected utility of f is at least as great as that of g – that is, P, which is to be interpreted as a relation of weak preference, is such that

$$(f, g) \in P \text{ iff } \sum_{i \in I} p(i)u(f(i)) \geq \sum_{i \in I} p(i)u(f(i)) \tag{2.2}$$

for every pair of alternatives f and g.

It is worth pausing to consider some of what is presupposed in (2.2), in particular those presuppositions that are not salient from the discussion of (2.1) and the expected utility approach to decision making under risk. First, it is presupposed that there are some reasonable ways of interpreting references to the *distinguished partition*, e.g, perhaps that it is maximally refined in some sense or maybe that it defined "locally" with reference to pragmatic considerations such as the decision maker's goals. These options in turn give rise to various questions of their own, e.g., questions concerning the existence of a maximally refined partition or even questions concerning the stability of the associated judgments of admissibility upon a not-necessarily-convergent series of increasingly refined partitions [36]. Second, an alternative, or *act*, such as f is required to be a function from the set of possible states of nature to the set of possible consequences. However, in (2.2), an argument to f is an element of the distinguished partition of the set of possible states, and of course such a thing is a subset of the set of possible states rather than an element of it. This apparent type-related predicament resolves to an innocent abuse of notation upon the assumption that f is invariant along each element of the partition in the sense that, for all i of the relevant partition and all states s, t in i, $f(s) = f(t)$.

Savage, as well as others, has offered axiomatic foundations for the extension of expected utility to decision making under uncertainty [22]. However, the normative status of such extensions has been called into question by examples of the sort made famous by Ellsberg [3, 4] and to some extent anticipated by Keynes [19] and Knight [20]. These examples purport to show instances of rational choice, in the pre-theoretical sense, which cannot be reconciled with Savage's axioms or any of the other familiar axiomatizations of subjective expected utility theory. Two sorts of axioms are of particular concern. Those of the first sort require that the agent satisfies optimization - i.e., the agent has a complete ranking of the alternatives and is disposed to select an available alternative that is optimal with respect to this ranking. Axioms of the second sort are independence assumptions, similar to what can be found throughout the general literature on additive models [21]. Roughly, in the context at issue, the relevant independence assumptions can be seen as requiring that states where f and g agree have no bearing on the agent's preferences concerning these acts, since if one of the states of agreement obtains then the outcome is a "sure thing" in the sense that it does not depend on whether f or g is chosen. These two types of conditions, ordering and independence, are of particular concern in relation to Ellsberg-type examples because such examples are supposed to demonstrate instances of rational choice that must violate at least one of them. While most of the theoretical options that have been offered in light of such examples maintain the assumption of optimization against a complete ordering [3, 5, 8] – and thus must relax the independence requirement in order to accommodate Ellsberg's example – others have argued that the best theoretical option is to abandon the idea that rationality is optimization against a complete ranking of the alternatives [25, 37, 46, 39]. Though demonstrably outside of the preference-based tradition of optimization – as they allow violations of the α and β conditions that were considered previously – theories of the sort offered by Levi [23, 25] can still be understood in terms of more general notions of admissibility as represented by set-valued choice functions, i.e., like the $\mathcal{C}$ component in the structures considered previously.

3

The previous section culminated with a brief discussion of some of the central difficulties that are associated with extending the expected utility hypothesis to decision making under uncertainty. In particular, difficulties associated with the Ellsberg examples can be used to motivate various alternatives to subjective expected utility theory. These alternative accounts differ from each other in significant ways – Levi's account abandons the classical framework of optimization in contrast to those accounts that accommodate Ellsberg-type examples by relaxing independence assumptions while still remaining within the classical, preference-based framework of optimization. Yet each of these alternative accounts that are cited at the end of the previous section exploit *indeterminate probabilities*, which are represented as sets of probability distributions. Early in the previous section it was observed that mere optimization is insufficient as a standard of individual rational choice since it does not require any connection between an agent's preferences, which are the things being optimized, and that agent's beliefs and desires. In the expected utility tradition, a rational agent's beliefs and desires are representable by a numerically precise probability function and a cardinal utility function, respectively. These two functions are the sole inputs to the decision rule of expected utility theory (2.2). Although Levi's account is unique among the main alternatives to subjective expected utility theory in that it abandons optimization, each of these alternative accounts can be seen as abandoning the Bayesian dogma that the rational agent's belief state can be represented by a numerically precise probability distribution. Despite significant differences between them, these alternatives to subjective expected utility theory each presuppose some antecedent judgment that is representable by a set of probability distributions. This raises the question as to whether these antecedent judgments can be motivated independently of the judgments of admissibility that they are supposed to inform. We now turn to one such motivation that is due to Levi but finds its roots in the pragmatist tradition of Peirce. Thus, we begin by drawing upon Olson's discussion in [30] and recall the relevant part of Peirce's philosophy, namely his *injunction against roadblocks in the path of inquiry*:

> Although it is better to be methodological in our investigations, and to consider the economics of research, yet there is no positive sin against logic in trying any theory which may come into our heads, so long as it is adapted in such a sense as to permit the investigation to go on unimpeded and undiscouraged . On the other hand, to set up a philosophy which barricades the road of further advance toward the truth is the one unpardonable offence in reasoning, as it is also the one to which metaphysicians have in all ages shown themselves the most addicted. (C.S Peirce)

Levi's argument in [23] against the Bayesian orthodoxy, in particular the Bayesian commitment to numerically precise probabilities, can be seen as rooted in his reconstruction of Peirce's injunction. Roughly, Levi's argument is that the Bayesian framework does not provide sufficiently neutral perspectives from which the agent can entertain rival theories, e.g., the various statistical hypotheses that are compatible with the information that the agent is given in Ellsberg-type scenarios. The introduction of indeterminate probabilities addressees this difficulty. Thus, confronted with the prospect of a single random selection from an urn that is known to be filled with exactly one hundred balls, which are identical to each other in all respects except that some are black and some are white and where nothing is known about the ratio of black balls to white balls in the urn, the good Bayesian's credal state concerning the two possible outcomes should be representable as a numerically precise probability. Thinking of the various candidate distributions as rival statistical hypotheses, our good Bayesian rules out all but one of these hypotheses, the one corresponding to its subjective probability, despite the fact that each of the candidate distributions is compatible with all that is known about the urn. In this sense the Bayesian orthodoxy urges a violation of Peirce's injunction – the strict Bayesian framework blocks paths that would allow the agent to *coherently* entertain statistical hypotheses that are incompatible with its subjective probability distribution. In contrast, the introduction of indeterminate probabilities allows for credal states – e.g., the one represented by the set of all candidate distributions – from which the agent can coherently entertain the entire range of statistical hypotheses that are consistent with what it knows about the urn.

In rejecting the requirement of numerically precise credal states, Levi, through his decision theory, rejects the classical standard of optimization as noted previously. In this way some departures from optimization can be seen as arising from a commitment to pragmatist philosophy. It should be noted that there are important motivations for abandoning optimization that are not clearly related to any essential element of pragmatism. For example, Sen's well-known examples concerning the "epistemic value of the menu" problem can be used to motivate rational violations of α and, thus, given the aformentioned necessity of α, a rejection of optimization. Yet these departures from optimization, much as optimization itself, can be accommodated within the wider framework of admissibility judgments, represented formally in terms of set-valued choice functions. In [9] I argue that such admissibility-based accounts are lacking as a foundation for rational choice and that *conditional* judgments of admissibility – represented formally in terms of poset-indexed families of the same sort of choice functions that are used to represent unconditional judgments of admissibility – provide a more appropriate basis for the study of rational choice. Whereas the rational agent's judgments of admissibility are supposed to be representable as a set-valued choice function $\mathcal{C}$ in the sense that $\mathcal{C}(X)$ is to be interpreted as the set of alternatives on menu X that the agent judges to be admissible, judgments of conditional admissibility are supposed to be representable as a conditional choice function χ in the sense that $\chi(e, X)$ is to be interpreted as the set of alternatives on menu X that the agent judges to be admissible *given* e, where e is among some partially-ordered set of potential "epistemic" states that are antecedent to judgments of admissibility. In [9] the suggested move from judgments of admissibility to conditional judgments of admissibility, like the Ellsberg-inspired move from preference to admissibility that is discussed at the end of the previous section, is apparently motivated only through a concern for rational choice. Yet the present section recalls a sense in which Levi's rejection of optimization can be seen as following from his interpretation of the Peircean injunction against placing roadblocks in the path of inquiry. This raises a question: *Can the move to conditional judgments of admissibility also be motivated through an appeal to Peircean pragmatism?* I believe that this question can be answered in the affirmative and I will attempt to sketch the relevant motivation in the remainder of the present article.

> The irritation of doubt causes a struggle to attain a state of belief. I shall term this struggle *Inquiry*, though it must be admitted that this is sometimes not a very apt designation.

> The irritation of doubt is the only immediate motive for the struggle to attain belief. It is certainly best for us that our beliefs should be such as may truly guide our actions so as to satisfy our desires; and this reflection will make us reject every belief which does not seem to have been so formed as to insure this result. But it will only do so by creating a doubt in the place of that belief. With the doubt, therefore, the struggle begins, and with the cessation of doubt it ends. (C.S. Peirce, from "The Fixation of Belief")

What does Peirce mean in suggesting that doubt is an irritant? At various places Peirce suggests that doubt is some sort of obstacle to action. The following passage from [27] supports such a reading:

> What is wrong with this state [doubt] is not that it is psychologically uncomfortable, but that it leads to paralysis of action. An inquirer has some end in view, and two different and inconsistent lines of action present themselves, bringing action to a halt ... (C.J. Misak from *Truth and the End of Inquiry: a Peircean Account of Truth*)

In addition to reminding us that Peirce's belief-doubt model of inquiry is intended as branch of logic, as opposed to a branch of psychology, Misak's remarks suggest that doubt, at least in the Peircean sense, interferes with some aspect of rational decision making.

We can consider at least two ways in which doubt might interfere with an admissibility-based account of rational decision making. Recall that according to the sort of account at issue a decision maker's state at time t determines a distinguished structure of type $\langle A, \mathcal{M}, \mathcal{C} \rangle$, where A is a nonempty set consisting of all things that could serve as alternatives, $\mathcal{M}$ is the set of all potential menus from A, and $\mathcal{C}$ is a function that associates each menu M with a set of admissible alternatives from M. One way in which doubt might interfere with an admissibility-based account is if its type-level requirements – i.e., that the decision maker's state at t determines a distinguished structure of a particular type – fail to cohere with the relevant notion of agency. More specifically, such interference would suggest a relevant notion of agency according to which a decision maker's state need not determine a distinguished structure of the appropriate type so long as that state is not free from doubt. This kind of interference will not be considered here, if for no other reason than that the type-level requirements of an admissibility account are modest relative to what is typically assumed of agents in theoretical work on rational choice. A second way in which doubt might interfere with an admissibility-based account of rational decision making is if its proposition-level requirements – i.e., conditions like α and β that can be construed as statements about structures that satisfy the aforementioned type-level requirements – fail to cohere with the relevant notion of agency. More specifically, such interference would suggest a relevant notion of agency according to which a decision maker's state always determines a distinguished structure of the appropriate type but, so long as that state is not free from doubt, the structure that it determines need not satisfy the relevant proposition-level requirements; e.g., in states that are not free from doubt the associated structure may violate α. Yet how according to this picture are the potential violations of the proposition-level requirements made salient to the inquiring agent itself as, presumably, must be the case if doubt is to irritate in the sense that Peirce requires? It seems that in order to make these potential violations salient to the inquiring agent itself one must assume that such an agent is capable of making conditional judgments of admissibility – in particular, it seems that such an agent must, at time t, be capable of judging admissibility on the supposition that it is in a certain state that is free from doubt, be capable of judging admissibility on the supposition that it is in a certain state that is not free from doubt, and be capable of recognizing the sense in which the proposition-level requirements of the given admissibility-based account are satisfied under the first supposition and violated under the second.

References

[1] M. Allais, *Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine*, Econometrica **21** (1953), 503–546.

[2] B. deFinetti, *La prévision: ses lois logiques, ses sources subjectives*, Annals de L'Institut Henri Poincaré **7** (1937), 1–68.

[3] D. Ellsberg, *Risk, ambiguity, and the Savage axioms*, The Quarterly Journal of Economics **75** (1961), 643–669.

[4] ______, *Risk, ambiguity and decision*, Garland Publishing, 2001.

[5] P. Gardenfors and N.E. Sahlin, *Unreliable probabilities, risk taking, and decision making*, Synthese **53** (1982), 361–386.

[6] G. Gigerenzer and R. Selten, *Bounded rationality: The adaptive toolbox*, MIT Press, 2001.

[7] I. Gilboa, *Rational choice*, MIT, 2010.

[8] I. Gilboa and D. Schmeidler, *Maximin expected utility with non-unique prior*, Journal of Mathematical Economics **18** (1989), 141–153.

[9] J. Helzner, *Rationalizing two-tiered choice functions through conditional choice*, Synthese (2011), In press.

[10] J. Helzner and V. Hendricks, *Agency and interaction: What we are and what we do in formal epistemology*, Journal of the Indian Council of Philosophical Research (2010).

[11] ______, *Formal epistemology*, Oxford Bibliographies Online, Oxford University Press, 2011.

[12] R. Jeffrey, *The logic of decision*, University of Chicago, 1983.

[13] J. Joyce, *The foundations of causal decision theory*, Cambridge, 1999.

[14] J. Kadane and P. Larkey, *Probability and the theory of games*, Management Science **28** (1982), no. 2.

[15] J. Kadane, M. Schervish, and T. Seidenfeld, *Rethinking the foundations of statistics*, Cambridge University, 1999.

[16] D. Kahneman and A. Tversky, *Choices, values, and frames*, Cambridge Univerity, 2000.

[17] M. Kaplan, *Decision theory as philosophy*, Cambridge Univerity, 1996.

[18] R. Keeney and H. Raiffa, *Decisions with multiple objectives: preferences and value tradeoffs*, Cambridge University, 1993.

[19] J.M. Keynes, *A treatise on probability*, MacMillan, 1921.

[20] F.H. Knight, *Risk, uncertainty and profit*, Houghton-Mifflin, 1921.

[21] D.H. Krantz, R.D. Luce, P. Suppes, and A. Tversky, *Foundations of measurment: Additive and polynomial representations*, Academic Press, 1971.

[22] D. Kreps, *Notes on the theory of choice*, Westview, 1988.

[23] I. Levi, *On indeterminate probabilities*, Journal of Philosophy **71** (1974), 391–418.

[24] ______, *The enterprise of knowledge*, MIT, 1980.

[25] ______, *The paradoxes of Allais and Ellsberg*, Economics and Philosophy **2** (1986), 23–53.

[26] R.D. Luce and H. Raiffa, *Games and decisions: Introduction and critical survey*, Dover Publications, 1989, The Dover edition is a republication of the 1957 work.

[27] C. Misak, *Truth and the end of inquiry*, Oxford, 1991.

[28] ______, *Isaac levi and his pragmatist lineage*, Essays on the Philosophy of Isaac Levi (E. Olson, ed.), Cambridge, 2006, pp. 18–31.

[29] L. Narens, *Abstract measurement theory*, MIT Press, 1985.

[30] E. Olson, *Isaac levi*, Essays on the Philosophy of Isaac Levi (E. Olson, ed.), Cambridge, 2006, pp. 18–31.

[31] C. Peirce, *The fixatation of belief*, Philosophical Writings of Peirce (J. Buchler, ed.), Dover, 1955, pp. 5–22.

[32] M. Peterson, *An introduction to decision theory*, Cambridge, 2009.

[33] F.P. Ramsey, *Truth and probability*, The Foundations of Mathematics and Other Logical Essays (R.B. Braithwaite, ed.), Routledge and Kegan Paul, 1931, pp. 156–198.

[34] M. Risse, *What is rational about nash equilibria?*, Synthese **124** (2000), no. 3.

[35] A. Rubinstein, *Lecture notes in microeconomic theory*, Princeton University Press, 2006.

[36] L.J. Savage, *The foundations of statistics*, Dover Publications, 1972, The Dover edition is a republication of the 1954 work.

[37] T. Seidenfeld, *Decision theory without 'independence' or without 'ordering': What is the difference?*, Economics and Philosophy **4** (1988), 267–290.

[38] A. Sen, *Choice functions and revealed preference*, Review of Economic Studies **38** (1971), no. 3, 307–317.

[39] ———, *Maximization and the act of choice*, Econometrica **65** (1997), no. 4, 745–779.

[40] ———, *Non-binary choice and preference*, Rationality and Freedom, Harvard, 2002, pp. 245–258.

[41] S.S. Stevens, *On the theory of scales of measurment*, Science **103** (1946), 677–680.

[42] P. Suppes, *Representation and invariance of scientific structures*, CSLI, 2008.

[43] M. Schervish T. Seidenfeld and J. Kadane, *Decisions without ordering*, Acting and Reflecting: The Interdisciplinary Turn in Philosophy (W. Sieg, ed.), Kluwer, 1990, pp. 143–170.

[44] B. van Fraassen, *Scientific representation*, Oxford, 2008.

[45] P. Wakker, *Prospect theory for risk and ambiguity*, Cambridge, 2010.

[46] P. Walley, *Statistical reasoning with imprecise probabilities*, Chapman and Hall, 1990.