

This is a section of [doi:10.7551/mitpress/11252.001.0001](https://doi.org/10.7551/mitpress/11252.001.0001)

The Handbook of Rationality

Edited by: Markus Knauff, Wolfgang Spohn

Citation:

The Handbook of Rationality

Edited by: Markus Knauff, Wolfgang Spohn


DOI: 10.7551/mitpress/11252.001.0001

ISBN (electronic): 9780262366175

Publisher: The MIT Press

Published: 2021

Funding for the open access edition was provided by the MIT Libraries Open Monograph Fund.



The MIT Press

8.4 Decision under Uncertainty

Brian Hill

Summary

A series of famous examples casts doubt on the standard, Bayesian account of belief and decision in situations of considerable uncertainty. They have spawned a significant literature in economics and, to a lesser extent, philosophy. This chapter surveys some of this literature, with an emphasis on the normative issue of rational decision.

1. The Challenge of Uncertainty

Imagine that you are faced with two urns, each containing only black and white balls. For one of the urns (the unknown urn), that is all you know; for the other (the known urn), you have counted the balls and know that exactly half are black. What would you say if asked for the probability that the next ball drawn from the known urn was black? And what about the unknown urn? And, if you had to place a bet on the next ball drawn from one of the urns being black, would you bet on the known or the unknown urn? And what if the bets were on the next ball being white?

This example crops up in several disciplines. Keynes (1921/2004) uses it to motivate his notion of *weight of evidence*: “It is evident that in either case the probability of drawing a white ball is $\frac{1}{2}$, but that the weight of the argument in favor of this conclusion is greater [for the known urn]” (p. 75). Replacing direct observation of the composition of the known urn with sampling yields something close to Popper’s “paradox of ideal evidence.” In economics, the example is associated with Ellsberg (1961) and has spawned a significant literature on *decision under uncertainty*¹ over the past 30 years.

Its importance lies in the challenge it poses to the standard account of rational belief and decision, Bayesianism (see chapter 8.2 by Peterson, this handbook), and in particular to

- the **Bayesian thesis about rational belief**: It can be represented by a function assigning a single number

(between 0 and 1) to each proposition or event, which satisfies the laws of probability;

- and the **Bayesian thesis about rational decision**: The chosen action in any decision is that which maximizes the *expected utility* or *desirability* on the basis of the agent’s beliefs.

Although not formulated as such, Keynes’s and Popper’s points appear to challenge the first thesis: representing belief by probabilities cannot, allegedly, capture the weight of evidence supporting a belief. As often noted, this argument is not watertight: there are several differences between one’s Bayesian probabilities concerning the two urns, including one’s beliefs about their composition or the robustness of beliefs to further observations (e.g., Joyce, 2005).

By contrast, Ellsberg’s version involves decision. He observed a tendency (borne out in subsequent experiments; Camerer & Weber, 1992) to prefer betting on the ball drawn from the known urn over betting on the one drawn from the unknown urn, whatever the color—a pattern of behavior that has come to be known as *uncertainty aversion* or *ambiguity aversion*. (The field often qualifies the unknown urn as involving *ambiguity* or *Knightian uncertainty*.²) Such preferences conflict with the Bayesian theses. A Bayesian would choose to bet on black from the known urn (call this event B_k) over black in the unknown urn (B_u) only if $p(B_k) > p(B_u)$, where p represents her beliefs. However, she would have the same preference over bets on white only if $p(W_k) > p(W_u)$. Clearly, no probability function p can satisfy both inequalities. So the Bayesian must condemn these so-called Ellsberg preferences as irrational. Some find this drastic. This opens up the question of what more permissive account should or can replace it.

Although distinct, the points concerning belief representation and decision are intimately related. The difference in your ignorance about the two urns—or in the “weight of evidence” in the two cases—seems to justify a preference between them (Ellsberg, 1961; Levi, 1986; Gilboa, Postlewaite, & Schmeidler, 2009). Indeed, Ellsberg’s

point about decision adds a twist to the debate about weight of evidence: whatever ways there are of capturing something like weight in the framework of Bayesian probability, they remain irrelevant for choice. Indeed, the two criticisms are perhaps strongest when combined: *Bayesianism*, it appears, *reserves no role for the weight of evidence (or similar factors) in choice*.

This survey will present some of the main responses to this challenge, particularly in the economic field of decision under uncertainty, but also in philosophy. While the focus will be on decision, the relationship with belief representation means that this issue cannot be ignored. Indeed, it will structure the survey: different belief representations (replacing Bayesian probabilities) naturally require different, although not unrelated, decision rules (replacing expected utility). The emphasis will be solely on rational decision, although the reader should bear in mind that some proposals were developed with other goals in mind (e.g., tractability for economic modeling or descriptive accuracy). For ease of exposition, we shall largely eschew technicalities, at times abstracting liberally from precise formulations to focus on the gist. For details, discussions of dimensions other than the normative one, or more on the relevant economic literature, the reader is referred to excellent and more complete existing surveys such as Gilboa and Marinacci (2013) and Machina and Siniscalchi (2014).

Table 8.4.1

Terminology and definitions

| Terminology | Explanation |
|--|---|
| S | The set of states of the world (state space) |
| X | The set of consequences (consequence space) |
| f, g , etc. | Acts (objects of choice): functions from S to X |
| E, F , etc. | Events: subsets of S |
| Δ | The set of probability measures over the state space S (probability space) |
| $EU_p f$ | The expected utility of act f calculated with probability measure $p \in \Delta$ and utility function u : $EU_p f = \sum_{s \in S} p(s)u(f(s))$ |
| $\mathbf{1}, \mathbf{0}, x, \dots$ | Consequence yielding utility value $\mathbf{1}, \mathbf{0}, x, \dots$ |
| $\mathbf{1}_E \mathbf{0}$ | Bet yielding consequence $\mathbf{1}$ if event E occurs and consequence $\mathbf{0}$ otherwise |
| The functional V represents the preference relation \succsim . | For every pair of acts f, g : $f \succsim g$ if and only if $V(f) \geq V(g)$. |

Table 8.4.1 lists the main terminology and definitions used throughout the chapter, which are standard in the economic branch of decision theory (although little depends on the use of this framework). For simplicity, we assume a fixed utility function throughout the chapter and use $\mathbf{1}, \mathbf{0}, x$ to refer to consequences yielding utility $\mathbf{1}, \mathbf{0}, x$, respectively, according to this function. We assume sets to be finite, as far as possible (so as to use sums instead of integrals). Some notions are illustrated on the previous example in table 8.4.2. The columns correspond to the states of the world, each specifying the resolution of all relevant uncertainty (the color of the ball drawn from each urn). The rows correspond to the acts—in this case, bets yielding 1 unit of utility if won and 0 otherwise. The entries in the table are the consequences obtained under each act and in each state (e.g., if you bet on B_k and the state $B_k \& B_u$ realizes, you obtain a consequence worth 1 utility unit).

2. Multiple Priors

A popular reaction to the opening example focuses on the *precision* of the Bayesian representation of beliefs. In the unknown urn, there is no evidence to justify a particular call on the probability of black on the next draw. Yet, by insisting that a precise probability value must be assigned to this event, Bayesianism has—the thought goes—no way of expressing this ignorance. A natural move is thus to allow imprecision in valuations, by using *sets* of probability measures as representations of belief. Bayesianism amounts to the special case where the agent's set is a singleton. Such a representation has been discussed and defended under names such as *credal sets* in philosophy (Levi, 1986; Joyce, 2010) and *imprecise probabilities* in statistics (Walley, 1991); in economics, one speaks of *multiple priors* or *sets of priors*.

How could or should one decide on the basis of a set \mathcal{C} of priors? One of the earliest theories in economics—*maximin Expected Utility* or *maximin-EU*, developed by Gilboa and Schmeidler (1989)³—looks at the worst-case expected utility across the set; that is, it considers the representation of preferences by

$$\min_{p \in \mathcal{C}} EU_p f. \quad (1)$$

The rule is *cautious* or *pessimistic*, insofar as it bases preferences on the worst the act can do, according to the probability measures in the set \mathcal{C} . As such, it can straightforwardly account for the Ellsberg preferences. (The interested reader may verify this using the set $\mathcal{C}^* = \{p \in \Delta: 0.3 \leq p(B_u) \leq 0.7; p(B_k) = 0.5\}$.)

Table 8.4.2
Ellsberg choices

| | B _k & B _u (Black ball drawn from both urns) | B _k & W _u (Black drawn from known urn, white from unknown urn) | W _k & B _u | W _k & W _u |
|---|---|--|---------------------------------|---------------------------------|
| 1 _{B_k} 0 (bet on B _k) | 1 | 1 | 0 | 0 |
| 1 _{B_u} 0 (bet on B _u) | 1 | 0 | 1 | 0 |
| 1 _{W_k} 0 (bet on W _k) | 0 | 0 | 1 | 1 |
| 1 _{W_u} 0 (bet on W _u) | 0 | 1 | 0 | 1 |

One purportedly restrictive aspect of this rule is the focus on the worst case; an apparently less extreme alternative is the α -maximin-EU or Hurwicz criterion, along the lines suggested by Hurwicz (1951; Jaffray, 1989):

$$\alpha \min_{p \in \mathcal{C}} EU_p f + (1 - \alpha) \max_{p \in \mathcal{C}} EU_p f, \tag{2}$$

where α is a number between 0 and 1. This rule contains maximin-EU as a special case (where $\alpha = 1$) but goes beyond the arguably extreme case of total caution by taking into consideration both the worst and the best case. Other generalizations in this direction allow the α to depend more or less strongly on the act f being evaluated (Ghirardato, Maccheroni, & Marinacci, 2004).

Another approach focuses on what an agent can say “for sure” on the basis of her set of priors. Under a popular proposal, a preference for one act over another is formed only if all probability measures in the set “agree” that the former has higher expected utility. In other words, for all acts $f, g, f \succcurlyeq g$ if and only if

$$EU_p f \geq EU_p g, \quad \text{for all } p \in \mathcal{C}. \tag{3}$$

This *unanimity rule* has been defended in economics (Bewley, 1986), in statistics under the name *maximality* (Walley, 1991), as well as by several philosophers (S. Bradley & Steele, 2016).⁴ Note that it does not order certain pairs of acts: the reader is invited to check, for instance, that it does not order the Ellsberg bets under the set of priors \mathcal{C}^* above. Related rules include Levi’s (1986) E-admissibility, which picks out acts that are best according to at least one probability measure in the set.

Beyond intuitive remarks about their reasonableness, how can decision rules such as these be evaluated on

normative grounds? Two potentially relevant families of results have been developed in the field; we shall illustrate them on maximin-EU.

2.1 Evaluation 1: Implications for Choice

One way of evaluating the normative credentials of a decision rule is in terms of its *implications for choice*. If it leads to unpalatable choices, such as choosing to obtain a sure loss, then this provides good reason for skepticism; if its consequences for choice seem sensible, then this may provide arguments in its favor. Any particular choice results from the combination of the decision rule (such as (1)) and the agent’s attitudes (the set of priors and utility function). To focus evaluation on the rule, decision theory considers the implications it has *no matter* the agent’s attitudes. For instance, the maximin-EU rule generates transitive preferences (see below) for any \mathcal{C} and u used. The central results in decision under uncertainty—*representation theorems*—fully characterize the implications of the decision rule for choice. That is, they provide necessary and sufficient conditions on preferences—called *axioms*—for there to be some specification of the agent’s attitudes that represents them according to the rule. Figure 8.4.1 presents the general schema on maximin-EU (for details, see Gilboa, 2009; Gilboa & Marinacci, 2013).

These results help pinpoint the properties of preferences that distinguish decision rules from one another. For instance, the Bayesian expected utility rule satisfies (see chapter 8.1 by Grüne-Yanoff and chapter 8.2 by Peterson, both in this handbook):⁵

Preferences \succcurlyeq satisfy a (particular) set of axioms if and only if there exists a set of priors \mathcal{C} and utility function u such that \succcurlyeq is represented by (1) calculated with \mathcal{C} and u .

Moreover, the representing \mathcal{C} and u are suitably unique.

Figure 8.4.1
Representation theorem.

Weak Order (WO): The preference relation is transitive (for all $f, g, h, f \succcurlyeq g$ and $g \succcurlyeq h$ imply $f \succcurlyeq h$) and complete (for all $f, g, f \succcurlyeq g$ or $g \succcurlyeq f$).

Sure Thing Principle (STP): The preference across any pair of acts is independent of what the acts yield on events where they agree.

For instance, in table 8.4.2, since $\mathbf{1}_{B_k}\mathbf{0}$ and $\mathbf{1}_{B_u}\mathbf{0}$ agree on the event that the balls drawn from the two urns are the same color (i.e., the first and fourth states), STP implies that preferences over these acts should be independent of their consequences on this event. The same goes for $\mathbf{1}_{W_k}\mathbf{0}$ and $\mathbf{1}_{W_u}\mathbf{0}$. However, since on the complement event (i.e., the event consisting of the second and third states), $\mathbf{1}_{B_k}\mathbf{0}$ coincides with $\mathbf{1}_{W_u}\mathbf{0}$ and $\mathbf{1}_{B_u}\mathbf{0}$ coincides with $\mathbf{1}_{W_k}\mathbf{0}$, STP implies that $\mathbf{1}_{B_k}\mathbf{0} \succcurlyeq \mathbf{1}_{B_u}\mathbf{0}$ if and only if $\mathbf{1}_{W_u}\mathbf{0} \succcurlyeq \mathbf{1}_{W_k}\mathbf{0}$. Hence, the Ellsberg preferences (for $\mathbf{1}_{B_k}\mathbf{0}$ over $\mathbf{1}_{B_u}\mathbf{0}$ and for $\mathbf{1}_{W_k}\mathbf{0}$ over $\mathbf{1}_{W_u}\mathbf{0}$) violate STP.

In terms of choice, the maximin-EU and unanimity rules take complementary approaches: the former weakens STP while retaining WO; the latter drops the completeness clause of WO while holding on to STP. Various arguments have been proposed, in both the philosophical and economic literatures, that violating one or other of these axioms leads to unsavory consequences, particularly in *dynamic* contexts. They have been used by some to argue for, say, dropping WO rather than STP (S. Bradley & Steele, 2016; Seidenfeld, 1988). They have also been used to criticize any divergence from Bayesianism—and hence all of the approaches discussed in this survey—as irrational (Hammond, 1988; Al-Najjar & Weinstein, 2009; Elga, 2010).⁶ This is currently the main battleground between Bayesianism and its critics and deserves a survey in itself. Note that even if such arguments are correct about the weaknesses of non-Bayesian approaches in dynamic contexts, these need to be traded off against their strengths in dealing with the role of evidence in choice. Some have claimed that the latter outweigh the former (Gilboa et al., 2009).

Apart from dropping STP, maximin-EU retains two axioms satisfied by expected utility:

P4: The preference for a bet on E over a bet on F with the same stakes is independent of the stakes.

Uncertainty Aversion (UA): For any pair of disjoint events E, F with $\mathbf{1}_E\mathbf{0} \sim \mathbf{1}_F\mathbf{0}$, $\frac{1}{2}\mathbf{1}_{E \cup F}\mathbf{0} \succcurlyeq \mathbf{1}_E\mathbf{0}$.

Preferences over bets on events—for a bet on B_k over a bet on B_u (table 8.4.2)—reflect agents’ “willingness to bet” and are typically considered to be related to their beliefs (Savage, 1954). P4 states that the relative willingness

to bet on different events is independent of the stakes involved—that is, of the consequences of winning or losing the bet, assuming that these are the same for both bets.⁷ Uncertainty Aversion translates the caution of maximin-EU. $\frac{1}{2}\mathbf{1}_{E \cup F}\mathbf{0}$ “hedges” the uncertainty involved in the bets on E and F : the uncertainty is “halved,” insofar as the payoff depends on whether E or F realizes, not which one does, but the winnings are halved as well. For an agent who dislikes uncertainty, such hedging can only be attractive, and hence the hedge is (weakly) preferred to the initial bets. This axiom is often taken as the property characterizing *uncertainty* (or *ambiguity*) *aversion*. However, it seems that people are not universally ambiguity averse (Wakker, 2010); for instance, in a version of the running example with 10 possible colors rather than 2, there is a tendency to bet on the unknown urn. Were such preferences to be deemed reasonable, this would be a blow for the rational credentials of maximin-EU. The α -maximin-EU rule retains P4 but drops UA and hence can accommodate these preferences.

2.2 Evaluation 2: Beliefs, Tastes, and Uncertainty Attitudes

Another criterion for evaluating a decision rule concerns its capacity to neatly separate beliefs from tastes or values. Consider the following oft-cited criticism of maximin-EU. In our running example, you know nothing about the composition of the unknown urn, so any probability measure is possible; the set $\mathcal{C}_0 = \{p \in \Delta: 0 \leq p(B_u) \leq 1, p(B_k) = 0.5\}$ captures this. However, according to maximin-EU with this set, a bet on the unknown urn is considered worse than a bet on black from an urn known to contain 1 black ball and 99 white ones. Such caution seems extreme.

This objection relies on two assumptions: first, that the set of priors involved in the decision rule represents beliefs and, second, that the beliefs should perfectly match available information. Concerning the latter assumption, the set \mathcal{C}_0 may reflect the objectively available information, but it does not capture a Bayesian agent’s beliefs, for these must be precise; so why should things be different for non-Bayesians? Indeed, the set of priors is often understood to represent the agent’s state of belief, which incorporates, but may go beyond, the “objective” information (see also R. Bradley, 2017). Indeed, many economists subscribe to the revealed-preference paradigm (see chapter 8.1 by Grüne-Yanoff, this handbook) and take the set of priors to be the uniquely determined set specified in the appropriate representation theorem (figure 8.4.1).

Turning to the first assumption, suppose that $\mathcal{C}_1 = \{p \in \Delta: 0.4 \leq p(B_0) \leq 0.6\}$ represents your preferences (according to (1)), so you do not hold the preference claimed in the objection.⁸ Since this set goes beyond the “objective” information, one can ask why you are using it. Is it because you have further information, or an inclination to “believe” in the principle of insufficient reason? Or is it instead because you are not so cautious as to use \mathcal{C}_0 but rather have a higher tolerance of uncertainty? These two possibilities are radically different: the former concerns beliefs (and their formation), whereas the latter relates to values or tastes for bearing uncertainty. Decision theory has developed formal tools that can shed light on such questions.

The central concept is comparative uncertainty aversion: between two agents, which (if any) is *more* uncertainty averse? The following is a widely accepted behavioral definition of the concept:

Comparative Uncertainty Aversion. Agent 1 is *more uncertainty averse* than agent 2 if, whenever agent 1 weakly prefers an act f over a sure c , then so does agent 2.

If Ann prefers an uncertain act f to a sure (utility) amount c , then she is not so averse to uncertainty as to consider bearing the uncertainty involved in f to be worse than getting the “non-uncertain” c . If Bob exhibits the same preference in all such cases, then he is not that averse to uncertainty either, so he is (weakly) less uncertainty averse than Ann.

Influenced by the analogy with risk attitude, economists have generally considered uncertainty attitude as a taste (for bearing uncertainty) and been interested in its consequences for, say, investment decisions. In particular, they have characterized comparative uncertainty aversion in terms of the parameters of decision models, with results such as that in figure 8.4.2 for maximin-EU (Ghirardato & Marinacci, 2002). It tells us that differences in uncertainty attitude correspond to differences in the set of priors. The fact that a comparison of tastes or values—as uncertainty attitudes are understood to be—translates to a difference in the set of priors casts doubt on the claim that this set reflects *only* beliefs. Indeed, the conclusion often drawn is that there is no clean interpretation of the set of priors: it reflects aspects of both belief and uncertainty attitude

(Klibanoff, Marinacci, & Mukerji, 2005). The maximin-EU account, it seems, does not have the resources to determine whether and to what extent the set \mathcal{C}_1 reflects enhanced beliefs or uncertainty tolerance.

These interpretational subtleties signal a perhaps more severe departure from the Bayesian benchmark than first imagined. Bayesianism vaunts a clear separation of beliefs (captured by the probability measure) and tastes (reflected in the utility function); maximin-EU apparently does not. To the extent that such a separation is desirable, this could bode ill for its normative credentials. For instance, in public decisions (e.g., concerning climate policy), it is customary for one group to supply the factual judgments (e.g., climate experts) and another to provide the values (e.g., society or its representatives); without the separation of beliefs and tastes, this division of labor is compromised. While current arguments against value-free science tend to focus on difficulties faced *in practice* (see Douglas, 2009, and chapter 14.2 by Bueter, this handbook), using sets of priors to report uncertainty risks ruling out such a possibility *in principle*. Moreover, the separation allows questions of theoretical rationality (e.g., learning) to be treated independently of those of practical rationality (e.g., decision; see chapter 2.2 by Wedgwood, this handbook). But if an agent’s set of priors reflects not only the evidence acquired but also her uncertainty attitude, then both will play a role in the formation of such sets; any theory of belief update will thus also have to incorporate value considerations related to uncertainty tolerance. To date, work on belief updating for sets of priors does not seem to have grappled with this issue.

2.3 Across Multiple Prior Rules

In comparing the various rules for choosing on the basis of sets of priors, the belief–taste separation issue seems not to favor any particular one but rather concerns the multiple prior representation in general: results à la figure 8.4.2 for the unanimity rule suggest that it suffers from similar problems.⁹ As concerns implications for choice, while some dispute the normative credentials of the preferences promoted by rules retaining WO (such as maximin-EU), few dispute the preference orderings yielded by the unanimity rule: if it recommends against

| | | |
|--|----------------|---|
| For agents 1 and 2 with the same utility function, agent 1 is more uncertainty averse than agent 2 | if and only if | agent 2’s set of priors is a subset of agent 1’s set of priors. |
|--|----------------|---|

Figure 8.4.2
Comparative Uncertainty Aversion.

an act as being worse than another for all probability measures, then it would seem like a bad idea to choose it. However, as noted previously, the rule may remain silent on some comparisons between acts; the issue is what to do in these cases.

Some suggest that there is nothing more to be said about rational decision: the unanimity rule provides all the guidance there is, and in cases where it is silent, there is no more guidance to be had (S. Bradley & Steele, 2016). Others invoke “mechanisms” that are specific to such cases of “indecision,” for instance, choosing a (contextually provided) status quo option, taking a deferral option, or choosing at random. Some have suggested “picking” a precise probability in the set of priors, or “sharpening,” for the purposes of decision (Joyce, 2010). Such a procedure needs to be carried out in a coherent way across decisions, to avoid agents making chains of decisions that yield sure losses. Rules such as maximin-EU can be thought of as providing principles for “picking” a probability: it always chooses a probability measure that evaluates the act as badly as possible among those in the set. Although the relevant probability measure differs according to the act evaluated (the measure used to evaluate a bet on white may not be relevant for a bet on black), the representation theorem tells us that the rule is invulnerable to the aforementioned problems. Indeed, each of the rules discussed above satisfying WO can be thought of as one way of “complementing” the unanimity approach in situations where it remains silent. The possibility of complementing approaches violating WO—seen as a strong rational base—by invoking considerations such as “security”—which may not enjoy the same interpersonally valid rational credentials—has been discussed by Levi (1986). Gilboa, Maccheroni, Marinacci, and Schmeidler (2010) invoke a similar intuition, distinguishing “objectively” from “subjectively rational” preferences, represented according to the unanimity and maximin-EU rules, respectively. They provide axioms characterizing the maximin-EU “completion” of given unanimity preferences.

3. Nonadditive Probabilities

Another reaction to Bayesianism’s apparent difficulties in accounting for evidence focuses on the additivity of probability functions. A range of proposals, including Dempster–Shafer belief functions, possibility functions, and Shackle’s degrees of surprise (see chapter 4.7 by Dubois & Prade, this handbook), employ real-valued functions assigning a number in $[0, 1]$ to each event and satisfying a mild monotonicity condition: bigger events

do not get a lower value. Such (monotonic real-valued set) functions are called *capacities* (other terms used, in various literatures, include “fuzzy,” “confidence,” and “plausibility measures”).

A first-pass decision rule involving capacities would keep the expected utility formula (table 8.4.1) but with capacities in the place of probabilities. It has long been known that such a rule has very unattractive behavioral properties: any decision maker using it will strictly prefer some dominated act (i.e., one that does worse in all states than another available act; see Quiggin, 1982; Wakker, 2010). Basically, the main decision rule under uncertainty using capacities that avoids such problems is the “Choquet Expected Utility” or “CEU” rule proposed by Schmeidler (1989), which evaluates an act f according to

$$\sum_{x_i} v(\{s: u(f(s)) \geq x_i\}) [x_i - x_{i+1}], \quad (4)$$

where v is the capacity, u is the utility function, and x_i are the utility values of consequences of f , in decreasing order. Beyond being the main rule suitable for the sorts of belief representation mentioned above, it has proved popular descriptively (see section 5 below), partly because of its implications for choice: the rule does not assume UA (although it does satisfy P4), which, as noted, is sometimes violated.

However, UA is satisfied whenever the capacity is convex, that is, whenever $v(E \cup F) + v(E \cap F) \geq v(E) + v(F)$ for all events E, F . In this case, the Choquet integral coincides with the maximin-EU evaluation using the set of probability measures dominating the capacity (called the *core*), that is, $\{p \in \Delta: \text{for all } E \subseteq S, p(E) \geq v(E)\}$. Since several nonadditive probability representations—such as Dempster–Shafer belief functions—are convex capacities, this means that the decision rule for them is (equivalent to) maximin-EU. So the previous remarks also hold for them. More generally, results à la figure 8.4.2 indicate that the capacity in the CEU rule reflects attitude to uncertainty (as well as, potentially, beliefs), so the separation issue remains problematic.

4. Confidence and Other Second-Order Representations

A third approach focuses on a purported multidimensional character of beliefs: to use Keynesian vocabulary, beyond the *balance* of evidence supporting a probabilistic judgment, there is also its *weight*. The thought that both of these “dimensions” are involved in the representation of beliefs dates back at least as far as Peirce (1878): “to express the proper state of our belief, not one number but two are requisite, the first depending on the inferred

probability, the second on the amount of knowledge on which that probability is based.” Similar distinctions have more recently played a prominent role in cognitive psychology (Griffin & Tversky, 1992). To capture this idea, representations generally employ some structure over the probability space Δ .

A simple example is a (weak) order over the probability space or, equivalently, a nested family of sets of probability measures, denoted Ξ (figure 8.4.3). R. Bradley (2017) and Hill (2013, 2019) argue that such a structure can capture an agent’s *confidence in her beliefs*. Each set in the family represents the beliefs or probability judgments held at a given level of confidence: the probability judgments that hold for all measures in the set are those that the agent holds with (at least) the corresponding amount of confidence. For instance, an agent represented according to figure 8.4.3 has high confidence in the probability judgment of 0.5 for B_k but only low confidence in the probability judgment of 0.5 for B_u . Larger sets in the family (for which fewer judgments hold) correspond to higher levels of confidence. Conceptually, confidence in a probability judgment is related to the evidence underpinning it (R. Bradley, 2017), so this approach relates to the aforementioned tradition. Technically, the representation is simply the “system of spheres” representation from belief revision and conditional logic (see chapter 5.2 by Rott and chapter 6.1 by Starr, both in this handbook), applied over the probability space rather than the state space. Belief representations in this spirit have been proposed by Gärdenfors and Sahlin (1982) and Nau (1992), although they assume that the second-order structure is cardinal: a number is assigned to each probability measure. Naturally, every cardinal structure over the probability space induces the ordinal structure described above (by ordering probability measures according to the value assigned). Similar uncertainty representations have also been promoted in econometrics (Manski, 2013).

Like multiple priors (section 2), several decision rules operate with this representation of beliefs. A notable family encapsulates the arguably reasonable maxim: *the higher the stakes involved in the decision, the more*

confidence is required in a belief for it to play a role. They involve a function D assigning to each decision a confidence level—formally, a set in the nested family Ξ representing the agent’s beliefs—and evaluate acts on the basis of the set picked out as appropriate for the decision by D . For each multiple-prior decision rule (section 2), there is a “confidence version,” involving the same rule but allowing the set of priors involved to vary according to the decision, and in particular the stakes involved in it. For instance, the maximin-EU member of this family represents preferences according to

$$\min_{p \in D(f)} EU_p f. \tag{5}$$

The main difference for choice with respect to maximin-EU lies in P4, which directly clashes with the intuition that, for different stakes, different levels of confidence may be appropriate and hence different beliefs may inform the decision, leading to potentially different willingness to bet. Decision rule (5) satisfies a weakening of this axiom, which allows the willingness to bet to change with the stakes, in line with the aforementioned maxim. Unlike maximin-EU, results à la figure 8.4.2 suggest a clean separation between beliefs and values (Hill, 2013): the family of sets Ξ represents the beliefs while the function D represents the agent’s uncertainty aversion, or taste for choosing on the basis of limited confidence. Drawing on similar results for the confidence version of the unanimity rule, Hill (2019) argues that these two points—the mild yet motivated divergence from multiple prior models in terms of choice implications, and the clean separation of beliefs and tastes—are general properties of the approach.

As noted previously, any cardinal representation of confidence in beliefs can be used with this sort of decision procedure, by “forgetting” the numbers and using only the order. However, other decision rules have been developed that make specific use of the cardinal structure. Prominent ones are the *variational preferences* rule (Maccheroni, Marinacci, & Rustichini, 2006), which represents preferences by

$$\min_{p \in \Delta} (EU_p f + c(p)), \tag{6}$$




Figure 8.4.3
Confidence in beliefs.

where c is a real-valued function on Δ , and the *confidence preferences* rule (Chateauneuf & Faro, 2009),

$$\min_{p \in \Delta} \frac{1}{\varphi(p)} EU_p f, \quad (7)$$

where φ is a function from Δ to $[0, 1]$. The former rule is motivated by an important literature on robustness in macroeconomics: a central model developed there (Hansen & Sargent, 2008) is a special case. The latter was motivated by and technically related to the literature on fuzzy sets: φ is a fuzzy set of probability measures. Both approaches diverge from maximin-EU by weakening P4, without motivating the divergence with a maxim similar to that behind (5). However, these models do not cleanly separate beliefs and tastes: results à la figure 8.4.2 suggest that c and φ capture uncertainty attitudes and hence cannot be thought of as pure representations of belief.

5. Probabilities after All

Finally, some strive to retain the probabilistic representation of beliefs. Inspired by work on decision under risk, one approach assumes that agents assign a precise probability $p(E)$ to each event E but that in choice, these probability values are “deformed” by a *weighting function* w , so that the decision weight attached to the event is $w(p(E))$. Since the composition of w and p is a capacity, this approach, popular in prospect theory (Wakker, 2010; see chapter 8.3 by Glöckner, this handbook), uses the CEU rule (section 3). A different approach reverts to a second-order probability μ over the space of first-order probabilities Δ . To the extent that this is a second-order structure, it can be thought of as related to the confidence approaches in section 4 (Marinacci, 2015).

The main challenge for these approaches lies in the incompatibility between the Ellsberg preferences and the property of *probabilistic sophistication*, which, roughly, states that all the decision-relevant information about an event is summarized in the probability assigned to it (Machina & Schmeidler, 1992). For instance, in the Ellsberg example, if the agent assigns probability 0.5 to both B_k and B_u , as is often deemed reasonable, then the deformed weight $w(0.5)$ assigned to the two events is the same, and the CEU rule predicts indifference between the bets. Likewise, a second-order probability μ generates a “reduced” probability $\Sigma p\mu(p)$, which can be used in the context of the expected utility rule; doing so is equivalent to the following evaluation, which applies expected utility at both levels:

$$\int_{\Delta} EU_p f d\mu(p) \quad (8)$$

However, as noted at the outset, expected utility cannot accommodate the Ellsberg preferences.

Faced with this situation, a common reply is to treat (the probabilities of) different events differently in decision according to the type of event, and more specifically the *source of uncertainty*, to which it belongs. The Ellsberg known and unknown urns are different sources of uncertainty, and so, the idea goes, the same probability assignment with respect to events from the different sources can be treated differently in decision.

Source dependence was introduced in the behavioral literature (Tversky & Fox, 1995) and continues to play a central role in current developments of one of the most important descriptive theories of decision, prospect theory (Abdellaoui, Baillon, Placido, & Wakker, 2011; Wakker, 2010; see chapter 8.3 by Glöckner, this handbook). Recent versions incorporate source preference by using the CEU rule with weighting functions that depend on the source of uncertainty, so probabilities about events concerning the known urn are weighted differently from those concerning the unknown urn.

For second-order probabilities, the following *smooth ambiguity* decision rule has been proposed (Klibanoff et al., 2005):¹⁰

$$\int_{\Delta} \varphi(EU_p f) d\mu(p), \quad (9)$$

with φ a real-valued function on utility values. The transformation function φ translates the difference in attitude to first-order and second-order uncertainty; interpreting the former as “physical uncertainty” and the latter as “model uncertainty,” it thus reflects different attitudes to these two sources of uncertainty (Marinacci, 2015). In particular, whenever φ is nonlinear, the attitudes differ, and the rule can accommodate Ellsberg preferences. This is one of the prominent models in the literature admitting a clean separation of beliefs from uncertainty attitudes: the second-order probability μ can be understood as a representation of the agent’s state of belief, whereas φ reflects her uncertainty attitude (Klibanoff et al., 2005). This separation, combined with the tractability and familiarity of the largely Bayesian framework, has contributed to its increasing popularity in economic modeling.

While some approaches invoking source dependence, particularly in the behavioral literature, have plainly descriptive aims, others seem to harbor normative ambitions (Marinacci, 2015) and hence call for an appraisal of their rational credentials. Criticisms have focused on their resolution of the challenge of reserving a role for weight of evidence in decision (section 1). They insist that the Bayesian representation captures all relevant

aspects of belief but introduce another element (the φ in (9)) to account for the role apparently played by weight of evidence. Some have suggested that this tension—in particular, the use of a parameter representing uncertainty attitude (φ) to reflect a characteristic of evidence, which is a factor pertaining to belief—is damaging for the rational credentials of approaches invoking source dependence (Hill, 2019).

Notes

1. We follow economists in using “decision under uncertainty” to refer to cases where the probabilities of the various outcomes are not provided, as opposed to “decision under risk,” where they are given.
2. After the notion of uncertainty introduced by Knight (1921).
3. Similar decision rules have been defended in philosophy (Gärdenfors & Sahlin, 1982) and developed in robust statistics, where they are called “ Γ -maximin” (Berger, 1985).
4. Different versions vary according to whether the inequalities are weak (as in (3)) or strict.
5. Throughout, we retain the standard names for axioms used in the literature. For further details, the reader is referred to Gilboa (2009) and Gilboa and Marinacci (2013), for example.
6. While these arguments focus on decision, they are related to “dilation” arguments that focus specifically on learning under imprecise probabilities (Seidenfeld & Wasserman, 1993; Joyce, 2010).
7. This axiom, which was introduced by Savage (1954), whose nomenclature we adopt, is related to Gilboa and Schmeidler’s C-Independence (1989).
8. Applying the maximin-EU rule with this set of priors evaluates a bet on black from the unknown urn using the worst-case probability for black—0.4—and this is better than the probability of getting black from an urn with one black ball out of 100.
9. For α -maximin-EU, despite some indications of belief–taste separation (Ghirardato, Maccheroni, & Marinacci, 2004), there are reasons for skepticism (Klibanoff, Marinacci, & Mukerji, 2005). It is complicated by the interdependence between the representing set of priors and the α (the uniqueness is weaker than in figure 8.4.1), which, without controversial assumptions, is uncondusive to a clean separation.
10. For related approaches, see Ergin and Gul (2009), Nau (2006), and Seo (2009).

References

Abdellaoui, M., Baillon, A., Placido, L., & Wakker, P. P. (2011). The rich domain of uncertainty: Source functions and their

experimental implementation. *American Economic Review*, 101(2), 695–723.

Al-Najjar, N. I., & Weinstein, J. (2009). The ambiguity aversion literature: A critical assessment. *Economics & Philosophy*, 25, 249–284.

Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*. New York, NY: Springer.

Bewley, T. F. (1986). Knightian decision theory. Part I. *Decisions in Economics and Finance*, 25(2), 79–110.

Bradley, R. (2017). *Decision theory with a human face*. Cambridge, England: Cambridge University Press.

Bradley, S., & Steele, K. (2016). Can free evidence be bad? Value of information for the imprecise probabilist. *Philosophy of Science*, 83(1), 1–28.

Camerer, C., & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5(4), 325–370.

Chateauneuf, A., & Faro, J. H. (2009). Ambiguity through confidence functions. *Journal of Mathematical Economics*, 45, 535–558.

Douglas, H. E. (2009). *Science, policy, and the value-free ideal*. Pittsburgh, PA: University of Pittsburgh Press.

Elga, A. (2010). Subjective probabilities should be sharp. *Philosopher’s Imprint*, 10. <http://hdl.handle.net/2027/spo.3521354.0010.005>

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

Ergin, H., & Gul, F. (2009). A theory of subjective compound lotteries. *Journal of Economic Theory*, 144(3), 899–929.

Gärdenfors, P., & Sahlin, N.-E. (1982). Unreliable probabilities, risk taking, and decision making. *Synthese*, 53(3), 361–386.

Ghirardato, P., Maccheroni, F., & Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118(2), 133–173.

Ghirardato, P., & Marinacci, M. (2002). Ambiguity made precise: A comparative foundation. *Journal of Economic Theory*, 102(2), 251–289.

Gilboa, I. (2009). *Theory of decision under uncertainty*. Cambridge, England: Cambridge University Press.

Gilboa, I., Maccheroni, F., Marinacci, M., & Schmeidler, D. (2010). Objective and subjective rationality in a multiple prior model. *Econometrica*, 78(2), 755–770.

Gilboa, I., & Marinacci, M. (2013). Ambiguity and the Bayesian paradigm. In D. Acemoglu, M. Arellano, & E. Dekel (Eds.), *Advances in economics and econometrics: Vol. 1. Economic theory: Tenth world congress* (Econometric Society Monographs, Vol. 49, pp. 179–242). Cambridge, England: Cambridge University Press.

- Gilboa, I., Postlewaite, A., & Schmeidler, D. (2009). Is it always rational to satisfy Savage's axioms? *Economics & Philosophy*, 25(3), 285–296.
- Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2), 141–153.
- Griffin, D., & Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24(3), 411–435.
- Hammond, P. J. (1988). Consequentialist foundations for expected utility theory. *Theory and Decision*, 25, 25–78.
- Hansen, L. P., & Sargent, T. J. (2008). *Robustness*. Princeton, NJ: Princeton University Press.
- Hill, B. (2013). Confidence and decision. *Games and Economic Behavior*, 82, 675–692.
- Hill, B. (2019). Confidence in beliefs and rational decision making. *Economics & Philosophy*, 35(2), 223–258.
- Hurwicz, L. (1951). Some specification problems and applications to econometric models. *Econometrica*, 19(3), 343–344.
- Jaffray, J.-Y. (1989). Linear utility theory for belief functions. *Operations Research Letters*, 8, 107–112.
- Joyce, J. M. (2005). How probabilities reflect evidence. *Philosophical Perspectives*, 19, 153–178.
- Joyce, J. M. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives*, 24, 281–323.
- Keynes, J. M. (2004). *A treatise on probability*. Mineola, NY: Dover. (Original work published 1921)
- Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6), 1849–1892.
- Knight, F. H. (1921). *Risk, uncertainty, and profit*. Boston, MA: Houghton Mifflin.
- Levi, I. (1986). *Hard choices: Decision making under unresolved conflict*. Cambridge, England: Cambridge University Press.
- Maccheroni, F., Marinacci, M., & Rustichini, A. (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6), 1447–1498.
- Machina, M. J., & Schmeidler, D. (1992). A more robust definition of subjective probability. *Econometrica*, 60(4), 745–780.
- Machina, M. J., & Siniscalchi, M. (2014). Ambiguity and ambiguity aversion. In M. J. Machina & W. K. Vicusi (Eds.), *Handbook of the economics of risk and uncertainty* (Vol. 1, pp. 729–807). Amsterdam, Netherlands: North-Holland.
- Manski, C. F. (2013). *Public policy in an uncertain world: Analysis and decisions*. Cambridge, MA: Harvard University Press.
- Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association*, 13(6), 1022–1100.
- Nau, R. F. (1992). Indeterminate probabilities on finite sets. *Annals of Statistics*, 20(4), 1737–1767.
- Nau, R. F. (2006). Uncertainty aversion with second-order utilities and probabilities. *Management Science*, 52, 136–145.
- Peirce, C. S. (1878). The probability of induction. *Popular Science Monthly*, 12(705), 83.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4), 323–343.
- Savage, L. J. (1954). *The foundations of statistics*. New York, NY: Dover.
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57, 571–587.
- Seidenfeld, T. (1988). Decision theory without “independence” or without “ordering.” *Economics & Philosophy*, 4, 267–290.
- Seidenfeld, T., & Wasserman, L. (1993). Dilation for sets of probabilities. *Annals of Statistics*, 21(3), 1139–1154.
- Seo, K. (2009). Ambiguity and second-order belief. *Econometrica*, 77, 1575–1605.
- Tversky, A., & Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review*, 102, 269–283.
- Wakker, P. P. (2010). *Prospect theory for risk and ambiguity*. Cambridge, England: Cambridge University Press.
- Walley, P. (1991). *Statistical reasoning with imprecise probabilities* (Monographs on Statistics and Applied Probability, Vol. 42). London, England: Chapman and Hall.

© 2021 The Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Knauff, Markus, editor. | Spohn, Wolfgang, editor.

Title: The handbook of rationality / edited by Markus Knauff and Wolfgang Spohn.

Description: Cambridge : The MIT Press, 2021. | Includes bibliographical references and index.

Identifiers: LCCN 2020048455 | ISBN 9780262045070 (hardcover)

Subjects: LCSH: Reasoning (Psychology) | Reason. | Cognitive psychology. | Logic. | Philosophy of mind.

Classification: LCC BF442 .H36 2021 | DDC 153.4/3—dc23

LC record available at <https://lcn.loc.gov/2020048455>