

Ambiguity and Ambiguity Aversion

Mark J. Machina^a, and Marciano Siniscalchi^b

^aDistinguished Professor of Economics, University of California, San Diego, CA, USA

^bDepartment of Economics, Northwestern University, Evanston, IL, USA

Contents

13.1 Introduction	730
13.2 Early Notions of Subjective Uncertainty and Ambiguity	732
13.3 The Classical Model of Subjective Probability	738
13.4 Ellsberg Urns	744
13.5 Models and Definitions of Ambiguity Aversion	750
13.6 Recent Definitions and Examples	777
13.7 Updating and Dynamic Choice	792
13.8 Conclusion	800
References	800

Abstract

The phenomena of *ambiguity* and *ambiguity aversion*, introduced in Daniel Ellsberg's seminal 1961 article, are ubiquitous in the real world and violate both the key rationality axioms and classic models of choice under uncertainty. In particular, they violate the hypothesis that individuals' uncertain beliefs can be represented by *subjective probabilities* (sometimes called *personal probabilities* or *priors*). This chapter begins with a review of early notions of subjective probability and Leonard Savage's joint axiomatic formalization of expected utility and subjective probability. It goes on to describe Ellsberg's classic urn paradoxes and the extensive experimental literature they have inspired. It continues with analytical descriptions of the numerous (primarily axiomatic) models of ambiguity aversion which have been developed by economic theorists, and concludes with a discussion of some current theoretical topics and newer examples of ambiguity aversion.

Keywords

Ambiguity, Ambiguity Aversion, Subjective Probability, Subjective Expected Utility, Ellsberg Paradox, Ellsberg Urns

JEL Classification Codes

D81 - Criteria for Decision-Making under Risk and Uncertainty, D80 - Information, Knowledge and Uncertainty, D01 - Microeconomic Behavior: Underlying

13.1 INTRODUCTION

Almost by its very nature, the phenomenon of uncertainty is ill-defined. Economists (and many others) agree that the uncertainty inherent in the flip of a fair coin, the uncertainty inherent in a one-shot horse race, and even the uncertainty inherent in the lack of knowledge of a deterministic fact (such as the 1,000,000th digit of π) are different notions, which may have different economic implications. Of the different forms of uncertainty, the phenomenon of *ambiguity*, and agents' attitudes toward it, is the most ill-defined.

The use of the term “ambiguity” to describe a particular type of uncertainty is due to Daniel Ellsberg in his classic 1961 article and 1962 PhD thesis,¹ who informally described it as:

“the nature of one’s information concerning the relative likelihood of events... a quality depending on the amount, type, reliability and ‘unanimity’ of information, and giving rise to one’s degree of ‘confidence’ in an estimation of relative likelihoods.” (1961, p.657)

As his primary examples, Ellsberg offered two thought-experiment decision problems, which remain the primary motivating factors of research on ambiguity and ambiguity aversion to the present day.² The most frequently cited of these, known as the *Three-Color Ellsberg Paradox*,³ consists of an urn containing 90 balls. Exactly 30 of these balls are known to be red, and each of the other 60 is either black or yellow, but the exact numbers of black versus yellow balls are unknown, and could be anywhere from 0:60 to 60:0. A ball will be drawn from the urn, and the decision maker is presented with two pairs of bets based on the color of the drawn ball.

	THREE-COLOR ELLSBERG PARADOX			TWO-URN ELLSBERG PARADOX			
	(single urn)			URN I		URN II	
	30 balls red	60 balls black yellow		100 balls red black		50 balls red	50 balls black
a_1	\$100	\$0	\$0	b_1	\$100	\$0	
a_2	\$0	\$100	\$0	b_2		\$100	\$0
a_3	\$100	\$0	\$100	b_3	\$0	\$100	
a_4	\$0	\$100	\$100	b_4		\$0	\$100

¹ Ellsberg (1961,1962). Ellsberg’s thesis has since been published as Ellsberg (2001).

² In (1961, p.653) and (1961, p.651, n.9) Ellsberg refers to Frank Knight’s (1921) “identical comparison” and to John Chipman’s (1958, 1960) “almost identical experiment” of the Two-Color Paradox, and in (1961, p.659, n.8) describes Nicholas Georgescu-Roegen’s (1954, 1958) notion of “credibility” as “a concept identical” to his own notion of ambiguity.

³ Ellsberg (1961, pp.653–656; 2001, pp. 155–158). Ellsberg (2001, pp. 137–142) discusses an essentially equivalent version with the payoffs \$100:\$0 replaced by $-\$100:\0 .

Ellsberg posited, and experimenters have confirmed,⁴ that decision makers would typically prefer bet a_1 over bet a_2 , and bet a_4 over bet a_3 , which can be termed *Ellsberg preferences* in this choice problem. Such preferences are termed “paradoxical” since they directly contradict the subjective probability hypothesis—if an individual did assign subjective probabilities to the events {red,black,yellow}, then the strict preference ranking $a_1 \succ a_2$ would reveal the strict subjective probability ranking $\text{prob}(\text{red}) > \text{prob}(\text{black})$, but the strict ranking $a_3 \prec a_4$ would reveal the strict ranking $\text{prob}(\text{red}) < \text{prob}(\text{black})$.

The widely accepted reason for these rankings is that while the bet a_1 guarantees a known probability 1/3 of winning the \$100 prize, the probability of winning offered by a_2 is unknown, and could be anywhere from 0 to 2/3. Although the range $[0, 2/3]$ has 1/3 as its midpoint, and there is no reason to expect any asymmetry, individuals seem to prefer the known to the unknown probability. Similarly, bet a_4 offers a guaranteed 2/3 chance of winning, whereas the probability offered by a_3 could be anywhere from 1/3 to 1. Again, individuals prefer the known-probability bet. Ellsberg described bets a_2 and a_3 as involving *ambiguity*, and a preference for known-probability over ambiguous bets is now known as *ambiguity aversion*.⁵

Ellsberg presented a second problem known as the *Two-Urn Paradox*, which posits a pair of urns, the first contains 100 black and red balls in unknown proportions, and the second contains exactly 50 black and 50 red balls.⁶ The decision maker is asked to rank the four bets shown in the Two-Urn Paradox table, where bet b_1 consists of drawing a ball from the first urn, and winning \$100 if it is black, etc. Agents are typically indifferent between b_1 and b_3 , and indifferent between b_2 and b_4 , but prefer the latter two bets over the former two, on the grounds that the latter two offer known probabilities of winning whereas the former two do not. Again, such preferences are incompatible with the existence of subjective probabilities—the ranking $b_1 \prec b_2$ would imply $\text{prob}(\text{red in Urn I}) < 1/2$, but the ranking $b_3 \prec b_4$ would imply $\text{prob}(\text{black in Urn I}) < 1/2$.

In this chapter we consider how economists have responded to these and similar examples of such “ambiguity averse” preferences. Section 13.2 gives an overview of early discussions of what has now come to be known as subjective uncertainty and the phenomenon of ambiguity. Section 13.3 reviews the classical approach to uncertainty, subjective probability and preferences over uncertain prospects. Section 13.4 presents the experimental and empirical evidence on attitudes toward ambiguity motivated by Ellsberg’s and similar examples. Section 13.5 gives analytical presentations of the most

⁴ See Section 13.4 of this chapter.

⁵ Individuals who would be indifferent between a_1 and a_2 , and between a_3 and a_4 , would be termed *ambiguity neutral*, and individuals who would prefer a_2 over a_1 and a_3 over a_4 would be termed *ambiguity loving*.

⁶ Ellsberg (1961, pp.650–651,653; 2001, pp.131–137). It is unfortunate that Fellner’s (1961) independent discovery, extensive discussion, and early experimental examination of the Two-Urn phenomenon has gone largely unrecognized.

important models of ambiguity and ambiguity aversion. Sections 13.6 and 13.7 present some recent developments in the field, and Section 13.8 concludes.

13.2 EARLY NOTIONS OF SUBJECTIVE UNCERTAINTY AND AMBIGUITY

13.2.1 Knight's Distinction

It is often asserted that the distinction between situations of probabilistic and nonprobabilistic beliefs was first made by Frank Knight (1921), in his use of the terms “risk” versus “uncertainty.” However, as LeRoy and Singell (1987) have convincingly demonstrated, Knight’s distinction between “risk” and “uncertainty” did *not* refer to the existence/absence of personal probabilistic beliefs, but rather, to the existence/absence of *objective probabilities* in the standard sense. In other words, Knight used “risk” to refer to situations where probabilities could either be theoretically deduced (“*a priori* probabilities”) or determined from empirical frequencies (“statistical probabilities”), and “uncertainty” to refer to situations that did not provide any such basis for objective probability measurement. However, Knight postulated that even under “uncertainty,” agents would still form *subjective* probabilities: “it is true, and the fact can hardly be over-emphasized, that a judgment of probability is actually made in such cases” (p.226) (Knight termed such probabilities “estimates”). Indeed, it is hard to find any more explicit adoption of the hypothesis of probabilistic sophistication under conditions of subjective uncertainty than Knight’s assertion that

“we must observe at the outset that when an individual instance [i.e., a one-time event] only is at issue, there is no difference for conduct between a measurable risk and an unmeasurable uncertainty. The individual, as already observed, throws his estimate of the value of an opinion into the probability form of ‘a successes in b trials’ (a/b being a proper fraction) and ‘feels’ toward it as toward any other probability situation.”⁷

Although Knight provided a verbal formulation of the concept of “subjective uncertainty,” the notion that agents in such situations might reject the standard probability calculus is due to his contemporary, John Maynard Keynes.

13.2.2 Keynes’ “Probabilities”

The fundamental concept in Keynes’ (1921) theory is a “probability,” which he defined as the “logical relation” between one proposition and another in situations where the

⁷ Knight (1921, p.234). Since he assumed that individuals always represented their beliefs by well-defined probabilities, what was the significance of the risk/uncertainty distinction for Knight? The answer is that under risk, probabilities are subject to independent measurement and hence are amenable to insurance, whereas under uncertainty they are not. Accordingly, only returns for bearing *uncertainty* should be attributed to a firm’s profits, since returns for bearing *risk* should be treated as costs (namely, the imputed cost of the firm’s decision to “self-insure” rather than purchase market insurance). See LeRoy and Singell (1987) for more complete discussion of this and other points.

first proposition neither logically assures nor logically excludes the second.⁸ For a given set of premises, therefore, the probability of a proposition is defined as the “rational degree of belief” that should be attached to it. Keynes did not consider “degree of belief” to be a personal or subjective notion, any more than its extreme cases of logical necessity or logical impossibility (say, of geometric propositions) are personal or subjective:

“The Theory of Probability is logical, therefore, because it is concerned with the degree of belief which it is rational to entertain in given conditions, and not merely with the actual beliefs of particular individuals, which may or may not be rational.” (p.4).

Keynes did allow some of his probabilities to take on numerical values, although

“the cases in which exact numerical measurement is possible are a very limited class, generally dependent on evidence which warrants a judgement of equiprobability by an application of the Principle of Indifference”⁹ (p.160).

Other probabilities, though not numerically measurable, can still be ranked:

“In these instances we can, perhaps, arrange the probabilities in an order of magnitude ... although there is no basis for an estimate how much stronger or weaker the [one probability] is than the [other]” (p.29).

However, some probabilities will not even be *ordinally* comparable:

“Is our expectation of rain, when we start out for a walk, always more likely than not, or less likely than not, or as likely as not? I am prepared to argue that on some of these occasions none of these alternatives hold...” (p.30).

Thus, a given pair of Keynesian probabilities can be related in one of three ways:

“I maintain ... that there are some pairs of probabilities between the members of which no comparison of magnitude is possible; that we can say, nevertheless, of some pairs of relations of probability that the one is greater and the other less, although it is not possible to measure the difference between them; and that in a very special type of case ... a meaning can be given to numerical comparisons of magnitude. I think that the results of observation, of which examples have been given earlier in this chapter, are consistent with this account.” (p.34).

In light of this, Keynes formally modeled his probabilities as a partial order (that is, transitive but not complete) with the following properties: all probabilities run from impossibility to certainty, certain subsets of probabilities form “ordered series” of mutually comparable elements, a probability may be a member of more than one ordered series, and all numerically measurable probabilities belong to a common ordered series. Since it allows for structures of belief which cannot be represented by numerical

⁸ Since propositions can take the form such as “the event *A* has occurred” and “the event *B* will occur,” this notion can also represent the relationship between a pair of events.

⁹ This was Keynes’ term for the Principle of Insufficient Reason (Section 13.2.5).

probabilities,¹⁰ Keynes' theory is the earliest example of a formal statement of non-probabilistic beliefs.¹¹

13.2.3 Shackle's "Potential Surprise"

The other early model of nonprobabilistic beliefs and preferences is that of George Shackle (1949a,b). The fundamental concept in Shackle's theory of belief is the "potential surprise" we would expect to experience upon learning that a particular event has occurred, or that a particular hypothesis is true. To distinguish this concept from standard probability, Shackle (1949a, p.113) gives the example of four equally qualified candidates for some appointment. A probabilistic representation of this situation may well assign each candidate a probability of 1/4, and hence view Candidate A's appointment as "unlikely." But given this symmetric uncertainty, we would hardly exhibit any "surprise" upon learning that Candidate A has received the position—nor, of course, would we be surprised to learn that it had gone to someone other than Candidate A. Moreover, these two surprise levels would remain at zero even if the number of equally qualified candidates rose from four to eight. On the other hand, if the pool were enlarged by the addition of clearly unqualified candidates, these new contenders would each be assigned a positive potential surprise.

The notion of potential surprise is geared toward the world *outside* of the gambling house, where our ignorance is not just in the relative likelihoods of a known set of alternatives, but in the very set of alternatives that might occur: "we need a measure of acceptance by which the individual can give to new rival hypotheses, which did not at first occur to him, some degree, and even the highest degree, of acceptance without reducing the degrees of acceptance accorded to any of those already in his mind" (1949b, p. 70)

Of course, by its very nature, such a measure of uncertainty will be nonadditive—as we have seen, the potential surprises of each member of an exhaustive set of events *could* all be zero, although according to Shackle, they *could not* all be positive. Shackle (1949a, App. E) gives their formal properties, including rules for their combination (e.g., the potential surprise of the union or the intersection of two events, of one event conditional upon another, etc.).

Just as his theory of beliefs departs from the traditional additive probability calculus, Shackle's theory of *preferences* over uncertain prospects departs from the additive expected value/expected utility approach. Consider an individual confronted with a set of alternative actions: "In order to assess the merits of any given course of action, a man must find some way of reducing the great array of hypotheses about the relevant consequences of this course ... to some compact and vivid statement" (1949a, p. 14).

¹⁰ Among other reasons, all likelihood relations represented by true numerical probabilities must be complete.

¹¹ See, however, Ramsey's (1926, §2) positive and normative criticisms of Keynes' theory.

To make this reduction, the individual will begin by determining the ability of each possible gain in an action to “stimulate him agreeably,” where this level of stimulation is an increasing function of the value of the gain and a decreasing function of its potential surprise. However, “the power of *mutually exclusive* hypotheses of success [alternative possible gains in a given action] to afford enjoyment by imagination is *not additive*” (1949a, p. 16). In fact, the entire positive stimulation of an action is defined to be that of its most stimulating possible gain:

“amongst all the hypotheses of success [potential gains] which the individual could entertain in regard to any venture, one alone is accountable in full for the enjoyment which he derives from the thought of this venture, and by itself determined the intensity of this enjoyment” (1949a, p. 16)

Similarly, the entire negative stimulation of an action is defined to be that of its most stimulating possible loss. Actions are then evaluated and ranked on the basis of “indifference maps” defined over such (stimulation of gain, stimulation of loss) pairs. Both the stimulation function and these indifference maps are amenable to theoretical analysis and empirical fitting, and Shackle applies his model to issues of gambling, investment, taxation and bargaining.

Although several writers¹² have criticized the unrealistic nature of some of his assumptions, Shackle’s work represents an admirable attempt to develop *and apply* a new mathematical theory of belief and decision under uncertainty, at a time when the expected utility model had not yet taken over the profession.

13.2.4 Ramsey’s “Degrees of Belief”

The earliest actual characterization of probabilistically sophisticated beliefs, in the sense of a set of assumptions on choice behavior which imply the existence of a classical probability measure over events, is that of Frank Ramsey (1926). Although he was probably not the first to observe that probabilistic beliefs could be measured by betting odds, he was the first to accomplish this without having to assume actual risk neutrality.

Since he was interested in the measurement of subjective probabilities—termed “degrees of belief”—rather than attitudes toward risk, Ramsey imposed the Bernoullian principle of expected utility maximization upon his agents, and indeed, worked directly in terms of the utilities, or as he called them, the “values,” of various outcomes. Accordingly, he assumed that “behaviour is governed by what is called the mathematical expectation [of utility or value]; that is to say, if \mathcal{P} is a proposition about which [the agent] is doubtful, any goods or bads for whose realization \mathcal{P} is in his view a necessary and sufficient condition enter into his calculations multiplied by the same fraction,

¹² E.g., Turvey (1949), Graaf and Baumol (1949), Carter (1950), Arrow (1951) and Ellsberg (1961).

which is called the ‘degree of his belief in \mathcal{P} .’ We thus define degree of belief in a way which presupposes the use of the mathematical expectation.” (1926, §3).

Ramsey defined a proposition \mathcal{P} to be *ethically neutral* if, holding all other aspects of the world constant, the individual is indifferent between its truth or falsity. The individual is said to have a “degree of belief $1/2$ ” in such a proposition if the prospects

$$\{\alpha \text{ if } \mathcal{P} \text{ is true; } \beta \text{ if } \mathcal{P} \text{ is false}\} \quad \text{and} \quad \{\beta \text{ if } \mathcal{P} \text{ is true; } \alpha \text{ if } \mathcal{P} \text{ is false}\}$$

are indifferent for all values α and β . Ramsey’s main assumptions are:

1. There exists at least one ethically neutral proposition with degree of belief $1/2$
2. If \mathcal{P} and \mathcal{Q} are both ethically neutral propositions with degree of belief $1/2$, and the individual is indifferent between the prospects

$$\{\alpha \text{ if } \mathcal{P} \text{ is true; } \beta \text{ if } \mathcal{P} \text{ is false}\} \quad \text{and} \quad \{\gamma \text{ if } \mathcal{P} \text{ is true; } \delta \text{ if } \mathcal{P} \text{ is false}\}$$

then he or she will be indifferent between the prospects

$$\{\alpha \text{ if } \mathcal{Q} \text{ is true; } \beta \text{ if } \mathcal{Q} \text{ is false}\} \quad \text{and} \quad \{\gamma \text{ if } \mathcal{Q} \text{ is true; } \delta \text{ if } \mathcal{Q} \text{ is false}\}$$

for all values α, β, γ and δ .

These assumptions, along with some technical ones, allowed Ramsey to identify the set of values with the real numbers, with the above preferences implying $\alpha + \beta = \gamma + \delta$. Having defined a way of measuring value/utility, he then invoked the principle of expectation to derive the individual’s beliefs, that is, their subjective probabilities of propositions or events: If the individual was indifferent between receiving α with certainty or the prospect $\{\beta \text{ if } \mathcal{R} \text{ is true; } \gamma \text{ if } \mathcal{R} \text{ is false}\}$, Ramsey defined their degree of belief in \mathcal{R} as $(\alpha - \gamma) / (\beta - \gamma)$, and assumed that this ratio would be the same for any other triple of values $\{\alpha', \beta', \gamma'\}$ that satisfy the same preference relation. Ramsey went on to derive notions such as the “conditional degree of belief in \mathcal{P} given \mathcal{Q} ” and to show that this concept of “degrees of belief” indeed satisfies the basic laws of probability theory.

Ramsey’s approach is limited (i) in that it imposes the property of expected utility maximization rather than jointly axiomatizes it, and (ii) in its dependence upon an essentially objective 50:50 randomization device (Assumption 1. above). (Both of these limitations are overcome by the approach of Savage (1954) described in Section 13.3.3.) However, since Ramsey was the first to characterize probabilistically sophisticated beliefs in terms of choice behavior, his insightful article deserves a prominent place in the literature.

13.2.5 Principle of Insufficient Reason

The earliest hypothesis concerning belief under subjective uncertainty is the so-called *Principle of Insufficient Reason*, which states that in situations where there is no logical or

empirical reason to favor any one of a set of mutually exclusive events or hypotheses over any other, we should assign them all equal probability.¹³ This principle is generally attributed to James (also known as “Jacob”) Bernoulli (1738). It was invoked by Bayes (1763) in his development of the binomial theorem (Stigler, 1986, pp.122–129) and by Laplace (1814) in his developments the Law of Succession and what is now called the Laplace distribution (Stigler, 1986, pp.109–113).¹⁴

Keynes (1921, Ch. IV), who also cites von Kries (1886), raised several objections to the Principle. The first relates to its implication that, in conditions of complete ignorance, we should assign equal probability to the validity of a hypothesis “this book is red” or to its complement. The problem of course is that the complement may consist of more than one mutually exclusive hypothesis (“this book is black,” “this book is blue,” etc.), and it is clearly impossible to assign a probability of 1/2 to *each* of these mutually exclusive hypotheses. A related objection also concerns multiple choice of partitions. If we have no information whatsoever as to the area or population of the regions of the world, then we would say that (i) a man is as likely to be an inhabitant of Great Britain as of France, and (ii) a man is as likely to be an inhabitant of *England* as of France. This, of course, would imply that Scotland and Wales are barren.

Another objection pertains to the application of the Principle to physical variables. Say we do know that the volume of a one-pound weight lies between 1 and 3 cubic inches, but have no further information on that value. This means that there is a 50:50 chance that its volume is greater than two cubic inches. On the other hand, our original information implies that the *density* of the object is between 1/3 and 1 pound/cubic inch, implying that there is a 50:50 chance that its density is greater than 2/3 pounds/cubic inch, which is inconsistent with the first conclusion.¹⁵

The most sophisticated of Keynes’ objections pertained to a situation identical to Urn I in Ellsberg’s Two-Urn example. In an urn with 100 black or red balls, does the Principle instruct us to treat all *ratios* of black to red balls (i.e., 0:100, 1:99, 2:98, ...) as equally likely, or does it instruct us to treat the color of each *individual* ball as equally likely to be black or red? Although the implications for betting on a single draw would be identical, the two conclusions have quite different implications for bets involving multiple draws.¹⁶

Formal axiomatic developments of the Principle of Insufficient Reason have been provided by researchers such as Chernoff (1954), Milnor (1954) and Sinn (1980).

¹³ Since this implies a uniform probability distribution over the events or hypotheses, it accordingly qualifies as a probabilistically sophisticated model of beliefs.

¹⁴ See also Keynes (1921, p.372) on Venn’s (1866) use of the Principle of Insufficient Reason in the Rule of Succession. On the other hand, see Shafer (1978) for arguments that at least some of James Bernoulli’s notions of “probability” were nonadditive.

¹⁵ A more economically-based example is that bond prices and interest rates cannot *both* have uniform probability densities.

¹⁶ See Savage (1954, pp. 63–67) for additional critical discussion of the Principle of Insufficient Reason.

13.3 THE CLASSICAL MODEL OF SUBJECTIVE PROBABILITY

13.3.1 Objective versus Subjective Uncertainty

Uncertain prospects can take different forms. A simple example of an *objectively uncertain prospect*—often called a *lottery* or a *roulette lottery*—is the gamble $\mathbf{P} = (x_1, p_1; \dots; x_n, p_n)$ yielding *outcome* x_i with a well-specified *objective probability* p_i . The outcomes in an objective lottery needn't be monetary; an objective lottery can be defined over any space of outcomes, such as standard consumption bundles, intertemporal time streams of monetary payments, vacations in different locales, etc. Nor need they be finite in number; the vector of probabilities (p_1, \dots, p_n) could be replaced by an arbitrary objective probability measure over outcome spaces in \mathbb{R}^1 or \mathbb{R}^n . The most general form of an objective lottery is that of an arbitrary probability measure $\pi(\cdot)$ over an arbitrary outcome space \mathcal{X} .

As mentioned, the uncertainty inherent in a fair coin or fair roulette wheel is distinct from the uncertainty inherent in a horse race or the weather. A *subjectively uncertain prospect*—often called an *act* or a *horse lottery*—is the bet $f(\cdot) = (x_1 \text{ if } E_1; \dots; x_n \text{ if } E_n)$ (or simply $(x_1, E_1; \dots; x_n, E_n)$) yielding x_j should the *event* E_j occur, for some mutually exclusive and exhaustive partition $\{E_1, \dots, E_n\}$ of all possible unfolding of the world, such as the partition $\{\text{horse 1 wins}, \dots, \text{horse } n \text{ wins}\}$. Partitions $\{E_1, \dots, E_n\}$ may in turn be thought of as alternative partitions (of varying coarseness) of an underlying space $\mathcal{S} = \{\dots, s, \dots\}$ of *states of nature*, which represents the subjective uncertainty at its finest and most basic level. Again, a subjective act needn't be finite-outcome; most generally, it consists of an arbitrary mapping $f(\cdot)$ from an arbitrary state space \mathcal{S} to an arbitrary outcome space \mathcal{X} . It is fair to say that, outside of gambling halls, most real-world uncertainty is subjective rather than objective.¹⁷

Uncertainty, be it objective or subjective, might well be resolved in two or more stages. A *two-stage* (or *compound*) *objective lottery* takes the form $(\dots; \mathbf{P}_i, p_i; \dots)$, yielding the objective lottery $\mathbf{P}_i = (\dots; x_{ik}, p_{ik}; \dots)$ with probability p_i , where for each i the probabilities (\dots, p_{ik}, \dots) sum to unity. A *two-stage subjective act* takes the form $(\dots; f_j(\cdot) \text{ if } E_j; \dots)$, yielding *subact* $f_j(\cdot) = (\dots; x_{jk} \text{ if } E_{jk}; \dots)$, where for each j the collection of *subevents* $\{\dots, E_{jk}, \dots\}$ is a partition of the event E_j . A two-stage *mixed* or *objective-subjective prospect*—termed a *horse-roulette act*, or sometimes an *Anscombe-Aumann act*—consists of a subjective act whose prizes are objective lotteries, and takes the form $(\dots; \mathbf{P}_j \text{ if } E_j; \dots) = (\dots; (\dots; x_{ij}, p_{ij}; \dots) \text{ if } E_j; \dots)$. Such prospects play an important role in the theory of ambiguity and ambiguity aversion. We analyze these in detail in Section 13.5.

Each two-stage objective lottery $(\dots; \mathbf{P}_i, p_i; \dots) = (\dots; (\dots; x_{ik}, p_{ik}; \dots), p_i; \dots)$ has a corresponding single-stage *reduced form lottery* $(\dots; x_{ik}, p_{ik} \cdot p_i; \dots)$, obtained by compounding

¹⁷ Although the first occurrence of this framework in its full generality seems to be Savage (1950) (in his review of Wald (1950)), it comes as a natural outgrowth of the statistical literature on hypothesis testing (Neyman and Pearson (1933), Wald (1939, 1950)), where the “states” were alternative hypotheses, “acts” were decisions to accept/reject the various hypotheses, and “consequences” were the (expected) values of the loss function.

the probabilities p_{ik} and p_i for each i, k .¹⁸ A decision maker may or may not be indifferent between a two-stage objective lottery and its corresponding reduced form—the hypothesis that they are in fact indifferent is known as the *Reduction of Compound Lotteries Axiom*.

Given a pair of objective lotteries $\mathbf{P} = (x_1, p_1; \dots; x_n, p_n)$ and $\mathbf{P}^* = (x_1^*, p_1^*; \dots; x_n^*, p_n^*)$ and some *mixture probability* $\alpha \in [0, 1]$, the $\alpha:(1-\alpha)$ *probability mixture* of \mathbf{P} and \mathbf{P}^* is the single-stage objective lottery $\alpha \cdot \mathbf{P} + (1-\alpha) \cdot \mathbf{P}^* = (x_1, \alpha \cdot p_1; \dots; x_n, \alpha \cdot p_n; x_1^*, (1-\alpha) \cdot p_1^*; \dots; x_n^*, (1-\alpha) \cdot p_n^*)$. The probability mixture $\alpha \cdot \mathbf{P} + (1-\alpha) \cdot \mathbf{P}^*$ of two lotteries is seen to be the single-stage reduced form of the two-stage compound lottery $(\mathbf{P}, \alpha; \mathbf{P}^*, (1-\alpha))$. A corresponding definition holds for probability mixtures $\alpha \cdot \pi(\cdot) + (1-\alpha) \cdot \pi^*(\cdot)$ of general objective lotteries.

Similarly, given two subjective acts $(x_1, E_1; \dots; x_n, E_n)$ and $(x_1^*, E_1; \dots; x_n^*, E_n)$ over a common partition¹⁹ $\{E_1, \dots, E_n\}$ of \mathcal{S} and a subset $\{E_1, \dots, E_m\}$ of these events, the $\{E_1, \dots, E_m\} : \{E_{m+1}, \dots, E_n\}$ *event mixture* of $(x_1, E_1; \dots; x_n, E_n)$ and $(x_1^*, E_1; \dots; x_n^*, E_n)$ is the (single-stage) act $(x_1, E_1; \dots; x_m, E_m; x_{m+1}^*, E_{m+1}; \dots; x_n^*, E_n)$ yielding outcome x_j if one of the events E_1, \dots, E_m should occur and x_j^* if one of E_{m+1}, \dots, E_n occurs. Given a pair of general subjective acts $f(\cdot)$ and $f^*(\cdot)$ and event $E \subseteq \mathcal{S}$, the $E : \sim E$ event mixture of $f(\cdot)$ and $f^*(\cdot)$ is the act $(f(\cdot), E; f^*(\cdot), \sim E)$ which yields outcome $f(s)$ for each state s in E and the outcome $f^*(s)$ for each state s in $\sim E$.

Although both objective lotteries and subjective acts can be defined more generally, from this point we restrict our attention to finite-outcome lotteries and acts.

13.3.2 Objective Expected Utility

The earliest and most basic model of preferences over uncertain prospects is the *objective expected utility* model, proposed by Bernoulli (1738) and formalized by von Neumann and Morgenstern (1944), Marschak (1950), Samuelson (1952) and others. In this model, preferences over objective lotteries can be represented by an ordinal *preference function* of the form $V(x_1, p_1; \dots; x_n, p_n) = \sum_{i=1}^n U(x_i) \cdot p_i$ or $V(\pi(\cdot)) = \int_{\mathcal{X}} U(x) \cdot d\pi(x)$, for some cardinal *von Neumann-Morgenstern utility function* $U(\cdot)$ over outcomes. Researchers such as Arrow (1963), Pratt (1964) and others have demonstrated how properties of the utility function $U(\cdot)$ correspond to features of attitudes toward objective uncertainty,²⁰ and the objective expected utility model has formed the cornerstone of the economic analysis of choice under uncertainty.

In addition to the usual properties corresponding to the existence of a preference ranking \succsim with a numerical representation $V(\cdot)$, the key feature of objective expected utility preferences, known as the *Independence Axiom*, is the property

¹⁸ Thus, the single-stage reduced form of the two-stage lottery $(\$10, 1/3; \$20, 2/3), 1/2; (\$0, 1/2; \$10, 1/6; \$30, 1/3), 1/2$ is the single-stage lottery $(\$0, 1/4; \$10, 1/4; \$20, 1/3; \$30, 1/6)$.

¹⁹ This partition could consist of any common refinement of the two acts' original partitions.

²⁰ See Chapter 3 of this Handbook.

Independence Axiom: For all lotteries $\mathbf{P}, \mathbf{P}^*, \hat{\mathbf{P}}$ and all $\alpha \in (0, 1]$, $\mathbf{P}^* \succcurlyeq \mathbf{P}$ if and only if $\alpha \cdot \mathbf{P}^* + (1 - \alpha) \cdot \hat{\mathbf{P}} \succcurlyeq \alpha \cdot \mathbf{P} + (1 - \alpha) \cdot \hat{\mathbf{P}}$.²¹

The intuition behind this property of preferences is most clearly revealed by thinking of the probability mixtures $\alpha \cdot \mathbf{P}^* + (1 - \alpha) \cdot \hat{\mathbf{P}}$ and $\alpha \cdot \mathbf{P} + (1 - \alpha) \cdot \hat{\mathbf{P}}$ in terms of their corresponding two-stage lotteries $(\mathbf{P}^*, \alpha; \hat{\mathbf{P}}, (1 - \alpha))$ and $(\mathbf{P}, \alpha; \hat{\mathbf{P}}, (1 - \alpha))$, where the first stage consists of the flip of a coin with objective probabilities $\alpha:(1 - \alpha)$ of landing heads:tails. Choosing between the two prospects essentially consists of choosing whether to receive \mathbf{P}^* or \mathbf{P} if it lands heads; if it lands tails two prospects will yield the same thing (namely $\hat{\mathbf{P}}$) anyway, so the decision maker should rank these two prospects in the same way he or she ranks \mathbf{P}^* and \mathbf{P} .

Paradoxes such as those of Allais (1953) have revealed systematic violations of the Independence Axiom, and have led to the development of *non-expected utility* models of preferences over objective lotteries. Such preferences are typically represented by functions $V(x_1, p_1; \dots; x_n, p_n)$, and several specific forms of such functions have been proposed.²²

13.3.3 Savage's Characterization of Subjective Expected Utility and Subjective Probability

As noted above, virtually all real-world uncertainty is subjective rather than objective, which led to the development of the corresponding *subjective expected utility* (SEU) model of Savage (1954). In this model, preferences over subjective acts are represented by an ordinal preference function of the form $W(x_1, E_1; \dots; x_n, E_n) = \sum_{j=1}^n U(x_j) \cdot \mu(E_j)$ or $W(f(\cdot)) = \int_{\mathcal{S}} U(f(s)) \cdot d\mu(s)$, for utility function $U(\cdot)$ and unique, additive *subjective probability measure*²³ $\mu(\cdot)$ over states. Just as the utility function $U(\cdot)$ represents an expected utility maximizer's attitudes toward risk, the subjective probability measure $\mu(\cdot)$ represents their beliefs of the likelihoods of the various states of nature and hence of the events based on them. Different decision makers can, and typically do, have different subjective probability measures (this, after all, is what makes for bets on horse races).

Savage obtained his characterization of subjective expected utility and subjective probability by means of the following axioms on a decision maker's preferences over subjective acts:²⁴

P1 Ordering: The preference relation \succcurlyeq is complete, reflexive and transitive.

²¹ Although implicitly invoked by von Neumann and Morgenstern in their formalization of the expected utility hypothesis (Malinvaud (1952)), the first formal statements of this property seem to be those of Marschak (1950) and Samuelson (1952).

²² See the Machina (1987) as well as Chapters 12 and 14 of this Handbook.

²³ Savage (1954) used the term *personal probabilities*.

²⁴ Axiom numbers are Savage's. Except for the Sure-Thing Principle, axiom names are our own. Savage (1954) provides an additional axiom, P7, used to extend his characterization to the case of infinite-outcome acts.

P2 Sure-Thing Principle: For all events E and all acts $f^*(\cdot), f(\cdot), f'(\cdot)$ and $f''(\cdot), (f^*(\cdot), E; f'(\cdot), \sim E) \succ (f(\cdot), E; f'(\cdot), \sim E)$ if and only if $(f^*(\cdot), E; f''(\cdot), \sim E) \succ (f(\cdot), E; f''(\cdot), \sim E)$.

P3 Eventwise Monotonicity: For all outcomes x^*, x , all nonnull²⁵ events E and all acts $f(\cdot), (x^*, E; f(\cdot), \sim E) \succ (x, E; f(\cdot), \sim E)$ if and only if $x^* \succ x$.

P4 Weak Comparative Probability: For all events E^*, E and all outcomes $x^* \succ x, \hat{x}^* \succ \hat{x}, (x^*, E^*; x, \sim E^*) \succ (x^*, E; x, \sim E)$ if and only if $(\hat{x}^*, E^*; \hat{x}, \sim E^*) \succ (\hat{x}^*, E; \hat{x}, \sim E)$.

P5 Non-Degeneracy: There exist outcomes x^* and x such that $x^* \succ x$.

P6 Small Event Continuity: For all acts $f^*(\cdot) \succ f(\cdot)$ and outcomes x , there exists a partition $\{E_1, \dots, E_n\}$ of \mathcal{S} such that both $f^*(\cdot) \succ (x, E_j; f(\cdot), \sim E_j)$ for all j and $(x, E_j; f^*(\cdot), \sim E_j) \succ f(\cdot)$ for all j .

Chapter 1 of this volume covers the axiomatic characterization of both objective and subjective expected utility. The axioms most relevant to the study of ambiguity and ambiguity aversion will turn out to be P2 (Sure-Thing Principle), P4 (Weak Comparative Probability) and a stronger version of P4 described in Section 13.3.5.

The intuition behind the Sure-Thing Principle is virtually identical to that of the Independence Axiom, with the coin replaced by an event E which may or may not occur. If two acts $(f^*(\cdot), E; f'(\cdot), \sim E)$ and $(f(\cdot), E; f'(\cdot), \sim E)$ yield the same outcome $f'(s)$ for each state s in the event $\sim E$, it should not matter what those statewise common outcomes are. Thus, replacing the common outcome $f'(s)$ with some different common outcome $f''(s)$ for each state s in $\sim E$ will not affect the preference ranking over the prospects. In the language of modern consumer theory, the Sure-Thing Principle states that preferences over subjective acts are *separable* across mutually exclusive events.

However, event-separability is only one of two distinguishing features of the subjective expected utility model, and by itself does not ensure the existence of well-defined subjective probabilities.²⁶ To P2 we must also add P4 (Weak Comparative Probability), which states that for any pair of events, the event on which the individual would prefer to stake the better of two prizes will not depend upon the prizes themselves. In other words, the decision maker has a well-defined *comparative likelihood ranking* over events. Together, the Sure-Thing Principle and Weak Comparative Probability Axiom form the heart of the subjective expected utility model.

13.3.4 Anscombe and Aumann’s Joint Objective-Subjective Approach

The contribution of Anscombe and Aumann’s (1963) joint objective-subjective approach is twofold. First, it provides a framework for representing uncertain prospects

²⁵ An event E is said to be *null* if, for all acts $f(\cdot)$ and $g(\cdot)$ such that $f(s) = g(s)$ for all $s \in \sim E$, it is the case that $f(\cdot) \sim g(\cdot)$ —that is, payoffs received on the event E do not matter.

²⁶ A decision maker with a “state-dependent” expected utility preference function $\int_s U(f(s)|s) \cdot d\mu(s)$ will be event-separable and hence satisfies the Sure-Thing Principle, but will not necessarily reveal well-defined likelihood rankings over events.

which involve both objective and subjective uncertainty. Such prospects play a key role in the field of ambiguity and ambiguity aversion—the key feature of Ellsberg urns is precisely that they involve both types of uncertainty. The second contribution is that by introducing objective prospects into the subjective framework, their approach allows for an axiomatic derivation of subjective probability which is considerably simpler than that of Savage (1954).

In addition to horse-roulette acts $(\dots; \mathbf{P}_j \text{ if } E_j; \dots) = (\dots; (\dots; x_{ij}, p_{ij}; \dots) \text{ if } E_j; \dots)$, Anscombe and Aumann consider three-stage compound prospects. These are objective lotteries $(\dots; f_k, p_k; \dots)$ whose “prizes” consist of horse-roulette acts $f_k(\cdot) = (\dots; \mathbf{P}_{jk} \text{ if } E_j; \dots) = (\dots; (\dots; x_{ijk}, p_{ijk}; \dots) \text{ if } E_{jk}; \dots)$. Such prospects can be termed *roulette-horse-roulette acts*. (A two-stage, horse-roulette act is thus a special case of a three-stage, roulette-horse-roulette act in which the first stage roulette lottery is degenerate.) It is important to note that, whereas horse-roulette acts $f_k(\cdot) = (\dots; \mathbf{P}_{jk} \text{ if } E_j; \dots)$ can involve different payoffs \dots, x_{ijk}, \dots , they are all defined over the same partition $\{\dots, E_j, \dots\}$. In other words, a roulette-horse-roulette act is a roulette wheel whose respective prizes are different bets on the same horse race (where the prizes can themselves be roulette lotteries). We note that, in the recent literature, the term “Anscombe–Aumann act” is usually reserved for two-stage, horse-roulette acts: see Section 13.5.

Because they are interested in deriving subjective probability, Anscombe and Aumann preassume that the individual has expected utility preferences over primitive objective lotteries $\mathbf{P} = (\dots; x_i, p_i; \dots)$, i.e., that there exists a von Neumann–Morgenstern utility function $U(\cdot)$ over final payoffs such that the expected utility of any primitive lottery is given by $U(\mathbf{P}) = \dots + p_i U(x_i) + \dots$. Furthermore, they assume that the individual also has expected-utility preferences over roulette-horse-roulette acts: there exists a von Neumann–Morgenstern utility function $W(\cdot)$ over horse-roulette acts such that the expected utility of the roulette-horse-roulette act $(\dots; f_k, p_k; \dots)$ is given by $\dots + p_k W(f_k) + \dots$. Note that the horse-roulette acts f_k are treated simply as prizes. Of course, these expected-utility preferences over primitive lotteries and roulette-horse-roulette acts satisfy the von Neumann–Morgenstern Independence axiom on the respective domains.

Their first assumption, which they term “Monotonicity in the Prizes,” is that if a pure roulette lottery $\hat{\mathbf{P}}_j$ is weakly preferred to \mathbf{P}_j , then the horse-roulette act $(\dots; \mathbf{P}_{j-1} \text{ if } E_{j-1}; \hat{\mathbf{P}}_j \text{ if } E_j; \mathbf{P}_{j+1} \text{ if } E_{j+1}; \dots)$ is weakly preferred to $(\dots; \mathbf{P}_{j-1} \text{ if } E_{j-1}; \mathbf{P}_j \text{ if } E_j; \mathbf{P}_{j+1} \text{ if } E_{j+1}; \dots)$, or in the author’s words, “if two horse lotteries are identical except for the prizes associated with one [horse], then your preference between the lotteries is governed by your preference between the prizes associated with that [horse].”

Their second assumption, “Reversal of Order in Compound Lotteries,” is that, for a given horse race $\{\dots, E_j, \dots\}$, given probability vector (\dots, p_k, \dots) and given collection of primitive objective lotteries $\{\mathbf{P}_{ij}\}_{i,j}$ to serve as prizes, the individual is indifferent between the roulette-horse-roulette acts

$$(\dots; (\dots; \mathbf{P}_{ij} \text{ if } E_j; \dots), p_i; \dots) \quad \text{and} \quad (\dots; (\dots; \mathbf{P}_{ij}, p_i; \dots) \text{ if } E_j; \dots)$$

(note that the first prospect is a nondegenerate roulette-horse-roulette act, but the second is actually a horse-roulette act). In Anscombe and Aumann’s words, “if the prize you receive is to be determined by both a horse race and the spin of a roulette wheel, then it is immaterial whether the wheel is spun before or after the race” (1963, p. 201).

These authors demonstrate how their mixed objective-subjective framework and set of assumptions imply the existence of well-defined subjective probabilities (\dots, q_j, \dots) over the states (in their setting, the horses), in the sense that the individual’s expected utility $W(f)$ of any horse-roulette act $f = (\dots; \mathbf{P}_j \text{ if } E_j; \dots)$ is given by $\dots + q_j \cdot U(\mathbf{P}_j) + \dots$ ²⁷

13.3.5 Probabilistic Sophistication

Although the Savage axioms imply both properties, it is possible for a decision maker to exhibit well-defined probabilistic beliefs without necessarily having expected utility risk preferences. A preference function $W(\cdot)$ over subjective acts is said to be *probabilistically sophisticated* (or satisfy the *Hypothesis of Probabilistic Sophistication*) if it takes the form $W(x_1, E_1; \dots; x_n, E_n) = V(x_1, \mu(E_1); \dots; x_n, \mu(E_n))$ for some probability measure $\mu(\cdot)$ over events and preference function $V(\cdot)$ over objective lotteries. Such a decision maker is accordingly indifferent between any subjective act $(x_1 \text{ on } E_1; \dots; x_n \text{ on } E_n)$ and its associated objective lottery $(x_1, \mu(E_1); \dots; x_n, \mu(E_n))$. Since the preference function $V(\cdot)$ needn’t take the expected utility form, such individuals are not necessarily subject to *Allais-type* violations of the Independence Axiom. However, since they retain the property of probabilistic beliefs, they are precisely the target of the *Ellsberg-type* effects—that is, the paradoxes of Section 13.1 and the additional effects reported in Section 13.4.

Machina and Schmeidler (1992, Thm.2)²⁸ have shown how the property of probabilistic sophistication can be characterized by dropping Savage’s Sure-Thing Principle P2 and strengthening his Weak Comparative Probability Axiom P4 to the following:

P4* Strong Comparative Probability: For all pairs of disjoint events E^*, E , all outcomes $x^* \succ x$ and $\hat{x}^* \succ \hat{x}$, and all acts $f^*(\cdot)$ and $f(\cdot)$, if $(x^*, E^*; x, E; f^*(\cdot), \sim(E^* \cup E)) \succ (x, E^*(\cdot); x^*, E; f^*(\cdot), \sim(E^* \cup E))$ then $(\hat{x}^*, E^*; \hat{x}, E; f(\cdot), \sim(E^* \cup E)) \succ (\hat{x}, E^*(\cdot); \hat{x}^*, E; f(\cdot), \sim(E^* \cup E))$.

Since the Sure-Thing Principle P2 and the Strong Comparative Probability Axiom P4* are independent properties of preferences, a decision maker could satisfy either one without the other—the probabilistically sophisticated form $W(x_1, E_1; \dots; x_n, E_n) = V(x_1, \mu(E_1); \dots; x_n, \mu(E_n))$ will satisfy P4* but generally not P2, whereas the

²⁷ Other joint axiomatizations of expected utility and subjective probability by means of an extraneous randomization device include those of Davidson and Suppes (1956), Pratt et al. (1964), DeGroot (1970, Ch.6), Fishburn (1970, Ch.6).

²⁸ See also Machina and Schmeidler (1995). Derivations of probabilistic sophistication under weaker assumptions have also been provided by Grant (1995) and Chew and Sagi (2006).

state-dependent expected utility form $W(\dots; x_j, s_j; \dots) = \sum_j U(x_j | s_j) \cdot \mu(s_j)$ will satisfy P2 but generally not P4*.²⁹

While most Ellsberg urn examples illustrate violations of both P2 and P4*, not all do, and it is departures from probabilistic sophistication (i.e., violations of P4*) which constitute the phenomena of ambiguity aversion or ambiguity preference.

13.4 ELLSBERG URNS

The examples of Section 13.1 are two of many proposed by Ellsberg and others of what have come to be known as *Ellsberg Urns*. Ellsberg's (1961) article contained another example, suggested to him by Kenneth Arrow, similar in spirit to the Two-Urn example but involving a single urn.³⁰ Again, the conjectured response is that decision makers would prefer the unambiguous bet c_1 over the ambiguous c_2 (which would imply $\text{prob}(\text{red}) > \text{prob}(\text{green})$), and prefer the unambiguous c_4 over the ambiguous c_3 (which would imply $\text{prob}(\text{red}) < \text{prob}(\text{green})$).

FOUR-COLOR ELLSBERG PARADOX				
(single urn)				
	100 balls		50 balls	50 balls
	black	red	green	yellow
c_1	\$100	\$100	\$0	\$0
c_2	\$100	\$0	\$100	\$0
c_3	\$0	\$100	\$0	\$100
c_4	\$0	\$0	\$100	\$100

13.4.1 Initial Reactions and Discussion

Because they struck at the heart of what many considered to be basic principles of rationality—the Sure-Thing Principle and probabilistic beliefs—Ellsberg's examples spawned a lot of discussion among the decision theory establishment. Ellsberg summarized some of their initial reactions as follows.³¹

"Responses do vary. There are those who do not violate the axioms, or say they won't, even in these situations (e.g., G. Debreu, R. Schlaifer, P. Samuelson); such subjects tend to apply the axioms rather than their intuition, and when in doubt, to apply some form of the Principle of Insufficient Reason. Some violate the axioms cheerfully, even with gusto (J. Marschak, N. Dalkey); others sadly but persistently.... Still others (H. Raiffa) tend, intuitively, to violate the axioms but feel guilty about it and go back into further analysis."

²⁹This is essentially the point of footnote 26.

³⁰ Ellsberg (1961, p.654, n.4). See also (1961, p.651, n.1).

³¹ Ellsberg (1961, pp. 655–656).

Ellsberg's report of a wide range of views on these issues played out in the subsequent literature. Whether or not he had any intuitive tendency for violation, Raiffa (1961) offered what has come to be the standard argument to an individual who would make the typical choices in the Three-Color Urn: If you really prefer a_1 over a_2 and a_4 over a_3 , then you presumably prefer a 50:50 coin flip of $a_1:a_4$ versus a 50:50 coin flip of $a_2:a_3$. But both coin flips reduce to a purely objective 50:50 coin flip of \$100:\$0. In Raiffa's view, "Something must give!" (1961, p.694)

Others have expressed a variety of views. Although Fellner (1961) largely supported a decision maker's right to possess Ellsberg-like preferences, he also asked whether a decision maker "is or is not likely gradually to lose this trait *as he gets used to the uncertainty with which he is faced*."³² Brewer (1963) argued that the rationality/irrationality of Ellsberg-type "slanting down of subjective probabilities" depends on whether or not a decision maker is allowed a free choice to bet on either an event or its complement—if not, he argues, Raiffa's comparison with a 50:50 coin flip won't apply (discussion of these issues continued in Fellner (1963) and Brewer and Fellner (1965)). Roberts (1963) reported, but did not accept, the argument that losing in either bet b_2 or bet b_4 on the 50:50 urn is somehow different from losing in either bet b_1 or bet b_3 on the unknown urn.³³ Smith (1969) and Sarin and Winkler (1992), however, suggest that a decision maker indeed does have distinct (and measurable) utility of money functions for prizes won from the two different urns. Finally, historians of economic thought are directed to the interesting 1961–1963 correspondence between Ellsberg and Leonard Savage (Savage (1963)).

13.4.2 Experiments on Ellsberg Urns and Ambiguity Aversion

Although Ellsberg himself only offered his examples as thought experiments, he recognized the need for formal experimentation from the very start.³⁴ The earliest reported experiments of this form seem to be those of Chipman (1958,1960).³⁵ Fellner (1961) offered various versions of the Two-Urn problem to a group of Yale undergraduates, and found an overall tendency to prefer the 50:50 rather than the unknown odds. Subsequent experiments by Becker and Brownson (1964), MacCrimmon (1968), Slovic and Tversky (1974), Curley and Yates (1989) and others also confirmed Ellsberg's conjecture of widespread ambiguity aversion. Although most of these experiments use students as subjects, researchers such as MacCrimmon (1965), Hogarth and Kunreuther (1989), Einhorn and Hogarth (1986), Viscusi and Chesson (1999), Ho et al. (2002) and

³² Fellner (1961, pp.678–679), original emphasis.

³³ "If I pick [bet b_1], I would be completely out of luck if there were no red balls in the urn." (p. 333). See also the discussion of Roberts (1963) and Ellsberg (1963) on Roberts' notion of "vagueness" in decision making and its implications for Ellsberg-type choice situations.

³⁴ "To test the predictive effectiveness of the axioms ... controlled experimentation is in order." (1961, p.655, n.6).

³⁵ See footnote 2.

Maffioletti and Santori (2005) have examined the ambiguity preferences of business owners, trade union leaders, actuaries, managers and executives, with the same overall findings.

In their own series of experiments, MacCrimmon and Larsson (1979) recognized that the interesting parameter in Ellsberg's examples was not the winning prize level, but rather, the amount of objective versus subjective uncertainty, which in the Three-Outcome Urn is given by the (known) proportion of red balls in the urn. MacCrimmon and Larsson's sequence of experiments accordingly set out to examine how subjects' choices depended on this proportion. In the standard specification of the Three-Color Urn, the proportion of red balls is $1/3$ (30 out of 90). Of course, if the proportion of red balls were actually zero, all subjects would prefer option a_2 in the first pair and a_4 in the second pair, which is consistent with the Sure-Thing Principle, and if this proportion were unity, all would now prefer a_1 in the first pair and a_3 in the second (also consistent with the Sure-Thing Principle). As the proportion increased from zero toward unity, the percentage of a_1 choices and a_3 choices should both rise, and under the hypothesis of ambiguity aversion, there would be some intermediate interval of probabilities within which a subject's choice would have flipped from a_2 to a_1 , but not yet flipped from a_4 to a_3 , yielding the classic Ellsberg-type violation of the Sure-Thing Principle for this urn. Using 100-ball urns, MacCrimmon and Larsson were able to present subjects with urns whose red-ball proportions took the values 0.20, 0.25, 0.30, 0.33, 0.34, 0.40, and 0.50. They indeed found such an intermediate interval of probabilities, and perhaps not surprisingly, the percentage of such violations was the greatest at $p = .33$.

Although each of Ellsberg's own examples pits a purely objective urn against an ambiguous urn with a fixed probability range,³⁶ researchers have also explored attitudes toward changes in the size of this range (holding the center constant). Becker and Brownson (1964), Larson (1980) and Viscusi and Magat (1992) did find an aversion to increases in the size of the range; Curley and Yates (1985) and Yates and Zukowski (1976) did not. In a reversal of Ellsberg's original specification of a fixed prize and ambiguous probability, Eliaz and Ortolevaz (2011) also found ambiguity aversion in the case of a fixed objective probability but an ambiguous prize level.³⁷ Du and Budescu (2005) found that subjects were willing to pay more to reduce the range ("vagueness") of outcome uncertainty than probability uncertainty.³⁸

While the phenomena of Reduction of Compound Lotteries and ambiguity neutrality are distinct properties, many researchers consider them to be closely related, and some of the rationality arguments against ambiguity aversion (such as Raiffa's (1961)) explicitly or implicitly invoke the reduction principle. In an experimental examination

³⁶ In the Three-Color Paradox the unknown probability of a black ball ranges from 0 to $2/3$; in the Two-Urn Paradox it ranges from 0 to 1; in the Four-Color Paradox it ranges from 0 to $1/2$.

³⁷ See these authors' further results in the case of joint outcome/probability ambiguity.

³⁸ Outcome and probability ranges were each classified into "low," "medium," or "high" levels of vagueness.

of this, Halevy (2007) appended two urns to the original two-color urns: one urn where the proportion of black balls satisfied a uniform objective distribution, and one urn which was either all black, or all red, with objective 50:50 probabilities. The two urns, together with Ellsberg's Urn I, are all purely objective, and each reduces to a 50:50 objective lottery over the best:worst monetary prize.³⁹ However, in Halevy's second urn all uncertainty is resolved in the first stage, whereas in Ellsberg's urn I all uncertainty is resolved in the second stage (Halevy's first urn involves both first stage and second stage uncertainty). Using this framework, Halevy found that subjects who satisfied the Reduction of Compound Lotteries Axiom under objective uncertainty were typically ambiguity neutral. Ozdenoren and Peck (2008) also took a dynamic approach, framing a two-stage Ellsberg urn problem as a "game against nature" and exploring various implications for ambiguity aversion and dynamic consistency.

Another systematic feature of attitudes toward ambiguity—reported in Ellsberg (1962, 2001) but not in Ellsberg (1961)—emerges from what he termed his *n-Color Example*.⁴⁰ Each of two urns contains 100 balls, from among 10 different colors. Urn I contains exactly 10 balls of each color, whereas Urn II contains the colors (and perhaps not all of them) in unknown quantities. Ellsberg naturally conjectured that most would prefer bet d_1 over d_3 , preferring d_1 's fixed 10% chance of getting \$0 to d_3 's unknown chance, which might be much higher. But he also felt that "a significant number" would *disprefer* bet d_2 to d_4 , dispreferring d_2 's fixed (but mere) 10% chance of winning \$100 to d_4 's unknown chance, which might also be much higher. As with the Two-Urn Paradox, the ranking $d_1 > d_3$ violates the hypothesis of probabilistic beliefs in the direction of ambiguity aversion.⁴¹ But by a similar argument, the ranking $d_2 < d_4$ violates probabilistic beliefs in the direction of *ambiguity preference*. In the language of Viscusi and Chesson (1999), the ambiguity in d_4 's small chance of winning the \$100 payoff allows for "hopes" that it might be much higher, whereas the ambiguity in d_3 's small chance of getting \$0 makes for "fears" that *it* might be much higher. A similar phenomenon occurs in bets involving losses: Kahn and Sarin (1988) found that subjects tended to be averse ("fearful") toward ambiguity in small chances of having to suffer a given loss, but prefer (be "hopeful" toward) ambiguity in small chances of not having to suffer it.⁴² Similar results were also obtained by Becker and Brownson (1964), Yates and Zukowski (1976), Curley and Yates (1985, 1989), Einhorn and Hogarth (1986), and Hogarth and Einhorn (1990). Cohen et al. (1985, 1987), however, found no correlation between subjects' *overall* levels of ambiguity aversion toward gains versus ambiguity aversion toward losses.

³⁹ Halevy used the prizes \$2:\$0 and \$20:\$0 for both his own and those of Ellsberg's urns used in his experiments.

⁴⁰ Ellsberg (1962, pp. 268–281; 2001, pp. 199–209). In his discussion, Ellsberg chooses $n = 10$, but feels that the effect would be more pronounced for $n = 100$.

⁴¹ $d_1 > d_3$ would imply $\text{prob}(\text{red in Urn II}) > 1/10$, which under the natural conjecture that $\text{prob}(\text{red in Urn II}) = \dots = \text{prob}(\text{mauve in Urn II})$, would imply $\text{prob}(\text{red in Urn II}) + \dots + \text{prob}(\text{mauve in Urn II}) > 1$. Similarly for $d_2 < d_4$.

⁴² Viscusi and Chesson (1999) found that for a potential loss, the threshold probability "at which ambiguity shifts from being a negatively valued fear to a positively valued hope" was about 1/2.

n-COLOR ELLSBERG PARADOX

		URN I									
		10 balls red	10 balls yellow	10 balls black	10 balls green	10 balls blue	10 balls purple	10 balls white	10 balls grey	10 balls orange	10 balls mauve
d_1		\$0	\$100	\$100	\$100	\$100	\$100	\$100	\$100	\$100	\$100
d_2		\$100	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0
		URN II									
		100 balls									
		red	yellow	black	green	blue	purple	white	grey	orange	mauve
d_3		\$0	\$100	\$100	\$100	\$100	\$100	\$100	\$100	\$100	\$100
d_4		\$100	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0

13.4.2.1 Forms of Preference Elicitation

One methodological issue, present throughout experimental work on choice and decision making, concerns exactly how preferences over Ellsberg-type prospects “should” be elicited. Ellsberg’s original presentations, and the great preponderance of subsequent experiments, simply presented subjects with pairs of alternatives, and asked for a direct choice within each pair. But standard consumer theory posits that the same ranking would be revealed if a subject’s preferences were instead assessed via an independent monetary valuation of each prospect,⁴³ with the valuations then compared. Fox and Tversky (1995), Chow and Sarin (2001) and Du and Budescu (2005) found that ambiguity aversion was reduced substantially (though not completely) when subjects were asked for separate monetary evaluations of ambiguous and unambiguous prospects (via their willingness to pay or willingness to accept) rather than asked for direct comparisons.

In another alternative to simple pairwise choice, MacCrimmon and Larsson (1979) presented subjects with sets of 11 prospects each, and asked them for a complete ranking of all prospects within each set (indifference was allowed). By including bets on stock index prices along with classic Ellsberg urns in their menus, these researchers were able to explore another question related to ambiguity, namely how subjects treated ambiguity in unknown urns with ambiguity outside of the laboratory.⁴⁴ MacCrimmon and Larsson found no net effect in either direction.

⁴³ Such as the procedure of Becker et al. (1964).

⁴⁴ Selten has suggested that many subjects may feel they can make “a very good estimate” of stock market events, and suggests investment in developing countries as better for such experiments (MacCrimmon (1968, p.28)).

13.4.2.2 Experimental Studies of Insurance and Medical Decisions Under Ambiguity

Other experimenters have also elicited subjects' ambiguity preferences in choices more relevant and realistic than simply drawing balls from urns. An obvious domain is that of insurance. Experiments on insurance decisions under ambiguity typically place subjects in the role of either consumers or suppliers of contracts such as flood or earthquake insurance, product warranties, etc. Although subjects are typically students, experiments and surveys by Einhorn and Hogarth (1986), Hogarth and Kunreuther (1992), Kunreuther (1989) and others have also found ambiguity aversion in hypothetical decisions by both professional actuaries and experienced insurance underwriters. Kunreuther et al. (1995) found ambiguity aversion in a field survey of primary-insurance underwriters in commercial property and casualty insurance companies. In experiments which included professional actuaries, and where subjects were asked to price insurance both as consumers and firms, Hogarth and Kunreuther (1989) found results which paralleled those of Kahn and Sarin (1988), Viscusi and Chesson (1999) and others as reported above, namely that both consumers and firms revealed ambiguity aversion toward low likelihood losses, which decreased as the likelihood of the loss increased. In an experiment involving real losses, Koch and Shunk (2013) found that ambiguity aversion was higher under unlimited liability than limited liability. Market and policy implication of ambiguity aversion are examined in Hogarth (1989), Camerer and Kunreuther (1989a,b), Hogarth and Kunreuther (1985) and Kunreuther and Hogarth (1992). Baillon et al. (2012) explore how different ambiguity attitudes play out in group belief aggregation.

Medical decisions by both patients and doctors, which also inherently involve ambiguity, have also been proposed to experimental subjects. Such experiments include decisions regarding vaccination of children (Ritov and Baron (1990)), heart disease (Curley et al. (1989)), residential location based on health risks (Viscusi et al. (1991), Viscusi and Magat (1992)) and others (e.g., Curley et al. (1984)). Gerrity et al. (1990) developed a multivariate measure of physicians' reactions to uncertainty, and used the results of an extensive survey to develop two "reliable and readily interpretable subscales" which they term "stress from uncertainty" and "reluctance to disclose uncertainty to others."

13.4.2.3 Additional Experiments on Ambiguity and Ambiguity Aversion

In other experiments involving real-world scenarios, Hogarth (1989) and Willham and Christensen-Szalanski (1993) gave subjects actual medical liability cases and manipulated ambiguity about the probability of winning in a legal scenario where hypothetical plaintiffs and defendants had to decide whether to go to court or settle out of court. In direct comparisons across contexts, Kahn and Sarin (1988) found that consumers' ambiguity attitudes differed across choices involving radio warranties, pharmaceutical

decisions and restaurant food quality. Maffioletti and Santori (2005) examined subjects' attitudes toward bets on real-world election results, and Baillon and Bleichrodt (2011) used bets on the temperature and on stock index prices. And while initially solely the realm of economists and psychologists, experimental work on decisions under ambiguity has now extended to the realm of neurology.⁴⁵

Ambiguity aversion has been and continues to be one of the most intensively experimentally explored phenomena in decision theory. Further discussion of this literature is provided in the surveys listed in Section 13.8.

13.5 MODELS AND DEFINITIONS OF AMBIGUITY AVERSION

Unlike the economic concepts of “risk” and “risk aversion,”⁴⁶ there is not unanimous agreement on what “ambiguity aversion,” or even “ambiguity” itself, exactly is. However several models and definitions have been proposed.

Most (though not all) of these models take as their starting point the following formalization of the objective/subjective uncertainty framework of Sections 13.3.1 and 13.3.4. Preferences are defined over the domain of horse-roulette acts—henceforth called *acts*—namely maps $f = (\dots; \mathbf{P}_j \text{ if } E_j; \dots) = (\dots; (\dots; x_{ij}, p_{ij}; \dots), E_j; \dots)$ from a (finite or infinite) state space \mathcal{S} to roulette lotteries \mathbf{P}_j over a set of prizes \mathcal{X} . The Independence property over this richer domain is identical to the Independence Axiom of objective expected utility, except for the more general notion of probability mixing it entails. Probability mixtures of horse-roulette acts are defined statewise: given acts $f = (\dots; \mathbf{P}_j \text{ if } E_j; \dots)$ and $g = (\dots; \mathbf{Q}_j \text{ if } E_j; \dots)$ over a common⁴⁷ partition $\{E_1, \dots, E_n\}$ of the state space \mathcal{S} , and probability $\alpha \in [0, 1]$, the mixture $\alpha \cdot f + (1 - \alpha) \cdot g$ is defined as the act

$$\alpha \cdot f + (1 - \alpha) \cdot g = (\dots; \alpha \cdot \mathbf{P}_j + (1 - \alpha) \cdot \mathbf{Q}_j; \dots)$$

The axioms that characterize subjective expected utility in this framework are accordingly⁴⁸

Weak Order: \succsim is complete and transitive.

Non-Degeneracy: There exist acts f and g for which $f \succ g$.

Continuity: For all acts f, g, h , if $f \succ g$ and $g \succ h$, there exist $\alpha, \beta \in (0, 1)$ such that $\alpha \cdot f + (1 - \alpha) \cdot h \succ g$ and $g \succ \beta \cdot f + (1 - \beta) \cdot h$.

Independence: For all acts f, g, h and all $\alpha \in (0, 1]$, $f \succsim g$ if and only if $\alpha \cdot f + (1 - \alpha) \cdot h \succsim \alpha \cdot g + (1 - \alpha) \cdot h$.

⁴⁵ See Hsu et al. (2005), Chew et al. (2008), Huettel et al. (2006), as well as the survey of Weber and Johnson (2008).

⁴⁶ E.g., Rothschild and Stiglitz (1970), Pratt (1964).

⁴⁷ As before, $\{E_1, \dots, E_n\}$ could be any common refinement of the two acts' original partitions.

⁴⁸ These versions of the expected utility axioms, due to Fishburn (1970), are referred to in the literature as the *Anscombe-Aumann axioms*. See also Schmeidler (1989).

Monotonicity: For all acts f, g , if the roulette lottery $f(s)$ is weakly preferred to the roulette lottery $g(s)$ for every state s , then $f \succcurlyeq g$.

The expected utility representation of preferences over horse-roulette acts $f = (\dots; \mathbf{P}_j$ if $E_j; \dots) = (\dots; (\dots; x_{ij}, p_{ij}; \dots), E_j; \dots)$ implied by these axioms takes the form

$$W(f) = \int_{\mathcal{S}} U(f(s)) d\mu(s) = \sum_{j=1}^n U(\mathbf{P}_j) \cdot \mu(E_j) = \sum_{j=1}^n \left[\sum_i U(x_{ij}) p_{ij} \right] \cdot \mu(E_j)$$

where $U(\cdot)$ is a von Neumann–Morgenstern utility function and μ is a finitely additive probability measure (“prior”), which is uniquely identified as in Savage’s axiomatization. As seen in the above equation, the term $U(f(s))$ in the integral $\int_{\mathcal{S}} U(f(s)) d\mu(s)$ is the *expected* utility of the roulette lottery $f(s)$. This is also the case for many of the models that we consider in this section, and that are axiomatized in the horse-roulette framework.

The above Independence axiom⁴⁹ turns out to imply the Sure-Thing Principle,⁵⁰ which implies that any Ellsberg-type violation of the Sure-Thing Principle is also a violation of Independence. It follows that any model of ambiguity aversion in the horse-roulette act framework must relax Independence.

Versions of the above axioms can also be stated in a setting closer to that of Savage, where acts are purely subjective horse lotteries. This requires that the set of prizes be suitably rich (for instance, an interval of the real line), with a suitable notion of “subjective mixture” on prizes. Nakamura (1990), Gul (1992) and Wakker (1989) take this approach, and Ghirardato et al. (2003) introduce a general notion of subjective mixture of prizes which allows a direct translation of the above axioms, and many of their relaxations are discussed in this section.

13.5.1 Maxmin Expected Utility / Expected Utility with Multiple-Priors

Gilboa and Schmeidler (1989, p.142) suggest the following explanation of the modal behavior in the Ellsberg Paradox:

“One conceivable explanation of this phenomenon which we adopt here is as follows: [...] the subject has too little information to form a prior. Hence (s)he considers a set of priors as possible. Being uncertainty averse, (s)he takes into account the minimal expected utility (over all priors in the set) while evaluating a bet.” (original emphasis)

⁴⁹ We distinguish between the two identically named conditions by the capitalization “Independence Axiom” for the Marschak/Samuelson axiom of Section 13.3.2 and “Independence axiom” for the current Independence property.

⁵⁰ Defining event mixtures as in Section 13.3.1, suppose $(f^*, E; f', \sim E) \succcurlyeq (f, E; f', \sim E)$. By Independence, $1/2 \cdot (f^*, E; f', \sim E) + 1/2 \cdot (f^*, E; f'', \sim E) \succcurlyeq 1/2 \cdot (f, E; f', \sim E) + 1/2 \cdot (f^*, E; f'', \sim E)$, which can be equivalently written as $1/2 \cdot (f^*, E; f'', \sim E) + 1/2 \cdot (f^*, E; f'', \sim E) \succcurlyeq 1/2 \cdot (f, E; f'', \sim E) + 1/2 \cdot (f^*, E; f', \sim E)$. Invoking Independence once again yields $(f^*, E; f'', \sim E) \succcurlyeq (f, E; f'', \sim E)$.

The resulting model is called *Maxmin Expected Utility (MEU)* or sometimes the *Multiple-Priors (MP)* model.⁵¹ Formally, consider a closed,⁵² convex set \mathcal{C} of probability measures—priors—on the state space \mathcal{S} , and a von Neumann–Morgenstern utility function $U(\cdot)$. An act $f(\cdot)$ is evaluated according to

$$W(f(\cdot)) = \min_{\mu \in \mathcal{C}} \int U(f(\cdot)) d\mu. \quad (13.1)$$

To see how this model allows for the typical preferences in Ellsberg's examples, consider the Three-Color Paradox of Section 13.1 (the analysis of the Two-Urn Paradox is similar). Let the state space be $\mathcal{S} = \{s_r, s_b, s_y\}$, where s_r denotes the draw of a red ball, etc., let the set of prizes be $\mathcal{X} = \{\$0, \$100\}$, and set $U(\$100) = 1$ and $U(\$0) = 0$. To reflect the assumption that 30 out of the 90 balls in the urn are red, but that the number of black and yellow balls is not known, consider the set of priors⁵³

$$\mathcal{C} = \{\mu \in \Delta(\mathcal{S}) : \mu(s_r) = 1/3\}. \quad (13.2)$$

Under this set of priors, the four acts in the Three-Color Paradox are evaluated as

$$W(a_1) = 1/3 \quad W(a_2) = 0 \quad W(a_3) = 1/3 \quad W(a_4) = 2/3$$

which implies the Ellsberg rankings $a_1 \succ a_2$ and $a_3 \prec a_4$.⁵⁴ To derive these values, observe that every prior $\mu \in \mathcal{C}$ assigns probability 1/3 to the state s_r , so that $W(a_1) = 1/3$. Similarly, every prior $\mu \in \mathcal{C}$ assigns probability 2/3 to the event $\{s_b, s_y\}$, so that $W(a_4) = 2/3$. Act a_2 yields \$100 on state s_b and \$0 otherwise; in other words, it is a bet on black. The prior in \mathcal{C} that assigns zero probability to s_b is the one that minimizes expected utility, and will accordingly be the one selected by the MEU criterion, so that $W(a_2) = 0$. In other words, the individual evaluates a bet on black as if *none* of the 60 unknown balls in the urn were black. Act a_3 yields \$100 on the event $\{s_r, s_y\}$ and zero otherwise; in other words, it is a bet *against* black. This time, the prior in \mathcal{C} that assigns unit probability to s_b (thus zero probability to s_y) is the one that minimizes expected utility and hence is the one selected, so that $W(a_3) = 1/3$. That is, the individual evaluates a bet *against* black as if *all* of the 60 unknown balls were black. This is a (stark) example of the “worst-case scenario” thinking embodied in equation (13.1). While the set \mathcal{C} used here is extreme, it is not the only one that generates the standard Ellsberg preferences: any set of priors where the probability

⁵¹ The expression “multiple-priors” is potentially ambiguous, because there are several well-known models which also employ sets of priors (for instance, the Variational and Smooth Ambiguity Preferences models).

⁵² If the state space is finite, the set \mathcal{C} is closed in the usual Euclidean topology. If it is infinite, it is closed in the weak* topology.

⁵³ $\Delta(\cdot)$ denotes the family of probability measures over a set. When no confusion can arise, to simplify notation, we write $\mu(\{s_r\})$, $\mu(\{s_r, s_b\})$, $\nu(\{s_r\})$, ... as $\mu(s_r)$, $\mu(s_r, s_b)$, $\nu(s_r)$, ... for measures (and later capacities) over singletons or finite sets.

⁵⁴ Note this also implies that a_3 is considered no better than a_1 , in spite of the fact that it yields a higher payoff on s_y .

of s_r is constant at $1/3$ and the probability of s_b ranges from less than $1/3$ to greater than $1/3$ will yield the above rankings.

As a historical note, Ellsberg himself proposed a decision criterion that is effectively a special case of *MEU*. He proposes that, by careful deliberation, an individual faced with an ambiguous situation may nevertheless “arrive at a composite ‘estimated’ distribution μ_0 that represents all his available information on relative likelihoods”; however, due to ambiguity, “[o]ut of the set $\Delta(\mathcal{S})$ of possible distributions there remains a set \mathcal{D} of distributions that still seem ‘reasonable,’ reflecting judgments that he ‘might almost as well’ have made, or that his information... does not permit him confidently to rule out.”⁵⁵ He then suggested (p.664) that individuals may evaluate acts according to the criterion

$$W(f(\cdot)) = \rho \cdot \int U(f(\cdot)) d\mu_0 + (1 - \rho) \cdot \min_{\mu \in \mathcal{D}} \int U(f(\cdot)) d\mu$$

where $\rho \in (0,1)$ represents the individual’s “degree of confidence” in the estimate μ_0 and $\mathcal{C} = \rho \cdot \mu_0 + (1 - \rho) \cdot \mathcal{D}$ is seen to be the set of priors. Kopylov (2006) analyzes this model when the set \mathcal{D} equals the set $\Delta(\mathcal{S})$ of all possible probability distributions—a specification which also appears in the literature on robust Bayesian analysis.

Gilboa and Schmeidler (1989) axiomatize the *MEU* decision criterion via axioms on horse-roulette acts. They retain the Weak Order, Monotonicity, Continuity and Non-Degeneracy axioms stated above, but weaken Independence, replacing it with

Certainty Independence: For all acts f, g , all constant acts x , and all $\alpha \in (0,1]$: $f \succcurlyeq g$ if and only if $\alpha \cdot f + (1-\alpha) \cdot x \succcurlyeq \alpha \cdot g + (1-\alpha) \cdot x$.

Uncertainty Aversion: For all acts f, g and all $\alpha \in (0,1]$: $f \succcurlyeq g$ implies $\alpha \cdot f + (1-\alpha) \cdot g \succcurlyeq g$.

That Independence must be relaxed follows from the fact that, as noted above, Independence implies Savage’s Postulate P2, and hence must be violated by Ellsberg-type preferences. The key question is to what extent Independence should be weakened. To gain some intuition, it is useful to add to the four acts a_1, \dots, a_4 of the Three-Color Paradox a fifth act a_5 representing a bet on yellow: specifically, a_5 yields \$100 if a yellow ball is drawn and \$0 otherwise. An individual with *MEU* preferences characterized by the set \mathcal{C} in equation (13.2) will be indifferent between betting on black or on yellow, that is, $a_2 \sim a_5$. However the mixture $1/2 \cdot a_2 + 1/2 \cdot a_5$ is strictly preferred to a_2 , which would not be possible with *EU* preferences. How should this be interpreted, and what is its relationship to ambiguity?

Act a_2 will yield \$100 if the ambiguous color black is drawn and \$0 if the ambiguous color yellow is drawn. Act a_5 instead yields \$100 if yellow and \$0 if black. Thus, whether

⁵⁵ Ellsberg (1961, p.661); notation in this paragraph adapted to the present chapter.

or not the better prize \$100 is obtained—if it is obtained at all—hinges crucially on which of the two ambiguous colors is drawn. (Both acts yield \$0 under the unambiguous color red.) By way of contrast, a 50:50 mixture of the two acts provides *hedging*: it removes the dependence of the prize on which of the two ambiguous colors is drawn. Specifically, the mixture $1/2 \cdot a_2 + 1/2 \cdot a_5$ yields the same objective roulette lottery (\$100, 1/2; \$0, 1/2) for *both* of the ambiguous colors yellow and black. The Uncertainty Aversion Axiom thus reflects a *preference for hedging*.⁵⁶ Mathematically, Uncertainty Aversion corresponds to *quasiconcavity* of the functional representation of preferences—an analytically convenient property.

On the other hand, mixing an act with a constant act (a constant prize or a constant objective lottery) does not provide such hedging, and there is less of an argument that such mixtures should necessarily be preferred (or dispreferred) by an ambiguity averter. The Certainty Independence Axiom accordingly requires that mixtures with *constant* acts do *not* affect preferences. While *MEU* preferences do satisfy Certainty Independence, there are reasons to relax or drop this axiom: see Sections 13.5.6–13.5.9.

Gilboa and Schmeidler (1989) show that the above axioms are necessary and sufficient for the existence of the *MEU* representation (13.1), with affine utility function $U(\cdot)$ and convex set of priors \mathcal{C} which is “identified up to convex closure”—that is, any other set of priors which represents preferences will have \mathcal{C} as its convex closure (the closure of its convex hull).

A generalization of the *MEU* model is the so-called α -maxmin, or α -*MEU* model

$$W(f(\cdot)) \equiv \alpha \cdot \min_{\mu \in \mathcal{C}} \int U(f(\cdot)) d\mu + (1 - \alpha) \cdot \max_{\mu \in \mathcal{C}} \int U(f(\cdot)) d\mu. \quad (13.3)$$

For $\alpha = 1$, this representation reduces to *MEU*, and for $\alpha = 0$ it reduces to what is termed *maxmax* expected utility $\max_{\mu \in \mathcal{C}} \int U(f(s)) d\mu$ and it allows for a whole range of intermediate attitudes toward ambiguity. Unfortunately, a general, behavioral characterization of this class of preferences is not available.⁵⁷ Furthermore, the intuitive interpretation of α as an ambiguity-aversion parameter is not warranted in general: Siniscalchi (2006) demonstrates that an α -*MEU* decision maker with $\alpha = 2/3$ and set of priors specified in (13.2) will have preferences in the Three-Color problem that are indistinguishable from one with $\alpha = 1$ (i.e., *MEU*) and set of priors $\{\mu \in \Delta(\mathcal{S}) : \mu(s_r) = 1/3, \mu(s_b) \geq 2/9\}$. The elements α and \mathcal{C} of the representation are not separately identified.

⁵⁶ The condition “ $f \succcurlyeq g$ implies $f \succcurlyeq \alpha \cdot f + (1 - \alpha) \cdot g$,” known as the *Uncertainty Loving* condition, captures the opposite intuition, and corresponds to the *maxmax* preference function defined below.

⁵⁷ Ghirardato et al. (2004) provide an axiomatization, but it turns out (Eichberger et al. (2011)) that their axioms can only hold for the cases $\alpha = 1$ and $\alpha = 0$. (Klibanoff et al. (2011)) provide a non-degenerate axiomatization of α -*MEU* when the state space describes the realizations of an exchangeable sequence of experiments.

Lehrer (2012) axiomatizes a special class of *MEU* preferences in which the set \mathcal{C} is generated by a partially specified probability. Axiomatizations of *MEU* which do not rely on horse-roulette acts, but require a rich outcome space, have been provided by Casadesus-Masanell et al. (2000a,b) and Ghirardato et al. (2003).

13.5.2 Choquet Expected Utility/Rank-Dependent Expected Utility

Another important model of ambiguity aversion, proposed by Schmeidler (1989), is the *Rank-Dependent* or *Choquet* model. Rather than capturing the individual’s perception of ambiguity by means of a family of priors as in the *MEU* model, it does so by means of a single *non-additive* probability measure (“*capacity*”) $\nu(\cdot)$. For example, in the Three-Color problem, the individual may have the beliefs

$$\nu(s_r) = 1/3, \quad \nu(s_b, s_y) = 2/3, \quad \nu(s_b) = \nu(s_y) = \varepsilon \tag{13.4}$$

for some $\varepsilon \in [0, 1/3)$. The fact that $\nu(s_b, s_y) > \nu(s_b) + \nu(s_y)$ indicates that there is ambiguity about the relative likelihood of a black vs. yellow draw: there is a residual probability mass equal to $\nu(s_b, s_y) - \nu(s_b) - \nu(s_y)$ that the individual, so to speak, does not know how to allocate between s_b and s_y . On the other hand, $\nu(s_b, s_y) + \nu(s_r) = 1 = \nu(\mathcal{S})$, which indicates that there is no ambiguity about the probability of drawing a red vs. a non-red ball. A nonadditive probability measure can encode both the individual’s assessment of relative likelihoods and the confidence attached to such assessment in a single function. Formally, a capacity is defined as a nonnegative real function $\nu(\cdot)$ defined over an algebra or σ -algebra of subsets of the state space \mathcal{S} that satisfies the normalization $\nu(\emptyset) = 0$ and $\nu(\mathcal{S}) = 1$, and the monotonicity property that $E \subseteq F$ implies $\nu(E) \leq \nu(F)$ (a standard probability measure is a capacity since additivity implies monotonicity). Capacities need not be additive, but they still retain the property that the weight assigned to a set is not smaller than the weight assigned to any subset.

Note that given this departure from additivity, specifying the capacity of the individual states in a finite state space does *not* fully characterize the capacity. For example, equations (13.4) do not pin down a unique capacity, because they leave $\nu(s_r, s_b)$ and $\nu(s_r, s_y)$ unspecified.

A key issue is how to compute the integral/expected value of a function with respect to a capacity. Consider a finitely-ranged function $g : \mathcal{S} \rightarrow \mathbb{R}$ that takes values $\alpha_1, \dots, \alpha_n$. The problem is that what might seem like the natural definition, namely

$$\int g(s) d\nu(s) = \sum_i \alpha_i \cdot \nu(\{s : g(s) = \alpha_i\})$$

can lead to failures of monotonicity, even in simple cases.⁵⁸

⁵⁸ For any capacity with $\nu(s_r) = \nu(s_b) = \nu(s_y) = 1/2$ and $\nu(s_r, s_b, s_y) = 1$, the act $\{7 \text{ if } s_r; 8 \text{ if } s_b; 9 \text{ if } s_y\}$ will be assigned the value $(1/2) \cdot 7 + (1/2) \cdot 8 + (1/2) \cdot 9 = 12$, which exceeds the value of 10 assigned to the dominating act $\{10 \text{ if } s_r; 10 \text{ if } s_b; 10 \text{ if } s_y\}$. A similar candidate, namely the formula $\sum_s g(s) \cdot \nu(s)$, is subject to the same difficulties.

To resolve this problem, Schmeidler (1986,1989) proposes the use of the *Choquet Integral*. Given a function $g : \mathcal{S} \rightarrow \mathbb{R}$ which takes values $\alpha_1 \geq \dots \geq \alpha_n$, the Choquet integral of g with respect to a capacity ν is defined as

$$\int g(s) d\nu \equiv \alpha_1 \cdot \nu(\{s : g(s) = \alpha_1\}) + \sum_{i=2}^n \alpha_i \cdot [\nu(\{s : g(s) \geq \alpha_i\}) - \nu(\{s : g(s) \geq \alpha_{i-1}\})]. \quad (13.5)$$

The telescoping property of the respective weights

$$\nu(\{s : g(s) = \alpha_1\}), \dots, [\nu(\{s : g(s) \geq \alpha_i\}) - \nu(\{s : g(s) \geq \alpha_{i-1}\})], \dots, [1 - \nu(\{s : g(s) \geq \alpha_{n-1}\})]$$

in this formula ensures that they sum to one, avoiding the difficulties noted in footnote 58. If the capacity ν is additive, and hence a probability measure, equation (13.5) reduces to the standard integral of g with respect to ν . For a general, bounded measurable function g , the Choquet integral is defined as

$$\int g d\nu = \int_{-\infty}^0 [\nu(\{s : g(s) \geq \alpha\}) - 1] d\alpha + \int_0^{\infty} \nu(\{s : g(s) \geq \alpha\}) d\alpha$$

where the integrals on the right side are in the sense of Riemann.

The *Choquet Expected Utility* (or *CEU*) representation of preferences over acts f is thus

$$W(f(\cdot)) \equiv \int U(f(\cdot)) d\nu \quad (13.6)$$

where $U(\cdot)$ is a von Neumann-Morgenstern utility, $\nu(\cdot)$ is a capacity, and the integral is in the sense of Choquet. This representation is sometimes also called *Rank-Dependent Expected Utility*, with reference to the analogous model of preferences under objective uncertainty.⁵⁹

To see that the *CEU* model can accommodate Ellsberg in the Three-Color example, consider the capacity ν defined by

$$\nu(s_r) = 1/3 \quad \nu(s_b, s_y) = 2/3 \quad \nu(s_b) = \nu(s_y) = 1/6 \quad \nu(s_r, s_b) = \nu(s_r, s_y) = 1/2$$

and set $U(\$100) = 1$ and $U(\$0) = 0$. Applying equation (13.5) yields $V(a_1) = 1/3$, $V(a_2) = 1/6$, $V(a_3) = 1/2$ and $V(a_4) = 2/3$, giving the Ellsberg rankings $a_1 \succ a_2$ and $a_3 \prec a_4$.

⁵⁹ E.g. Quiggin (1982), Yaari (1987), Segal (1987b). See also Section 13.5.3 below.

CEU preferences capture ambiguity aversion (or preference) by allowing the individual to exhibit a preference (or aversion) for *hedging*. Observe that a 50:50 probability mixture of the ambiguous acts b_1 and b_3 in the Two-Color problem yields the completely objective (and to an ambiguity averter, strictly preferred) lottery $(\$100; 1/2; \$0, 1/2)$. This occurs, of course, because b_1 yields its best prize \$100 for a red draw, whereas b_3 yields its best prize for a black draw—the two acts are, in a sense, “negatively correlated,” so that mixing them in this way reduces the variability of outcomes (or more generally, utility) across states. Given a pair of multiple-state acts $f = [x_1 \text{ if } s_1; \dots; x_n \text{ if } s_n]$ and $g = [\gamma_1 \text{ if } s_1; \dots; \gamma_n \text{ if } s_n]$, this effect will occur, in part, for any a pair of states s_i, s_j for which $x_i > x_j$ and $\gamma_i < \gamma_j$. When two acts are *comonotonic*, that is, when $x_i \geq x_j \Leftrightarrow \gamma_i \geq \gamma_j$ for all i, j , probability mixing cannot lead to any hedging.

The Choquet model captures ambiguity aversion/preference by weakening the Sure-Thing Principle just enough to allow such a preference/aversion for hedging—that is, preferences are allowed to violate the key mixture property of the axiom so long as, but *only* so long as, there is actually some potential for hedging. Whenever there is *no* such potential, that is, whenever the acts being mixed are comonotonic, the mixture property must continue to hold:

Comonotonic Independence: For every triple of pairwise comonotonic acts f, g, h , and every $\alpha \in (0, 1]: f \succcurlyeq g$ if and only if $\alpha \cdot f + (1-\alpha) \cdot h \succcurlyeq \alpha \cdot g + (1-\alpha) \cdot h$.

Schmeidler (1986) shows that a function defined on simple (or bounded) measurable functions is monotonic, as well as both monotonic and additive with respect to comonotonic functions, if and only if it is a Choquet integral. Building on this, Schmeidler (1989) shows that *CEU* preferences are characterized by Weak Order, Continuity, Monotonicity, Non-Degeneracy, and the Comonotonic Independence Axiom. The uniqueness properties of *CEU* are the same as those for *SEU*: the capacity ν is unique, and utility is unique up to a positive affine transformation.

As noted, *CEU* preferences needn't be uncertainty-averse. Schmeidler (1989) demonstrates that a *CEU* preference function will satisfy the Uncertainty Aversion Axiom of Section 13.5.1 if and only if its capacity ν is *convex*: that is, if for every pair of events E, F , $\nu(E \cup F) + \nu(E \cap F) \geq \nu(E) + \nu(F)$. Uncertainty-averse *CEU* preferences are thus also a (strict) special case of *MEU* preferences, where the set \mathcal{C} of priors coincides with the *core* of the capacity ν :

$$\mathcal{C} = \text{core}(\nu) = \{\mu \in \Delta(\mathcal{S}) : \forall E, \mu(E) \geq \nu(E)\}.$$

Gilboa (1987) and Sarin and Wakker (1992) provide axiomatizations of *CEU* in the original Savage framework. Nakamura (1990) provides another, in a setting where the state space can be finite, provided the set of prizes is suitably rich. Another axiomatization of *CEU* in a fully subjective setting with rich outcomes can be found in Ghirardato et al. (2003). The most complete exposition of the rank-dependent model is that of

Wakker (2010). Lehrer (2009) proposes an integral for capacities that coincides with the Choquet integral if the capacity is convex, but is different in general.

13.5.3 Segal's Recursive Model

The *Recursive Model* of Segal (1987a) is the earliest of a group of models which can be described as “two-stage.” Ambiguity is modeled by assuming that the probability distribution μ over the state space \mathcal{S} is not known, but random with probability measure M over $\Delta(\mathcal{S})$. In the first stage, one particular probability measure $\mu \in \Delta(\mathcal{S})$ is realized. In the second stage, the state s is drawn according to the distribution μ . The first stage may be only an idealization, a fictitious construct that is helpful to describe a particular way to evaluate acts; the second stage is real. Here we assume that all measures have finite-support.

Under this model, an individual treats a basic act $(\dots; x_j \text{ if } E_j; \dots)$ as the two-stage prospect $(\dots; (\dots; x_j, E_j; \dots), \mu_k; \dots)$, where μ_k , the probability distribution (prior) over states, is the realization of the first-stage uncertainty, and E_j , the realized event under distribution μ_k , is the realization of second stage uncertainty.⁶⁰ Each such act has a corresponding two-stage objective lottery of the form $(\dots; (\dots; x_j, \mu^{-1}(x_j); \dots), M(\mu_k); \dots)$. The key question is how to evaluate such two-stage acts and lotteries.

An individual who applies the Reduction of Compound Lotteries (ROCL) axiom of Section 13.3.1 to collapse this two-stage objective lottery to a single-stage one cannot exhibit Ellsberg-type preferences: such an individual will be probabilistically sophisticated, with statewise subjective probabilities $\sum_k \mu_k(s_i) \cdot M(\mu_k)$. Segal, however, argues that it is appropriate to drop ROCL in such a setting:

“...if a sufficiently long time passes between the two stages of the lottery, then there is no reason to make this reduction assumption... It is my belief that decision makers consider the Ellsberg urn as a real two-stage lottery, in which the first, imaginary stage and the second, real stage are clearly distinguishable. Therefore, they do not feel themselves obliged to obey the reduction of compound lotteries axiom...” (1987a, p.178)

Segal begins by assuming the individual has a preference function $V(\cdot)$ defined over single-stage lotteries. Given an act $f(\cdot) = (\dots; x_j, E_j; \dots)$, each measure μ over \mathcal{S} induces a simple lottery of the form $(\dots; x_j, \mu(E_j); \dots)$ and the individual uses $V(\cdot)$ to determine the *certainty equivalent* $CE(f, \mu)$ of this lottery, that is, the value for which

$$V(CE(f, \mu), 1) = V(\dots; x_j, \mu(E_j); \dots).$$

The individual's uncertainty over the measure $\mu(\cdot)$, as represented by the measure $M(\cdot)$ over the space $\Delta(\mathcal{S})$, generates the real-valued lottery $(\dots; CE(f, \mu_k), M(\mu_k); \dots)$.⁶¹

⁶⁰ Since it involves the orthogonal state spaces \mathcal{S} and $\Delta(\mathcal{S})$, the mixture $(\dots; (\dots; x_j, E_j; \dots), \mu_k; \dots)$ is to be distinguished from a *two-stage subjective act* $(\dots; f_j(\cdot) \text{ if } E_j; \dots)$ as defined in Section 13.3.1, which involves subpartitions of $(\dots; \{\dots, E_{jk}, \dots\}; \dots)$ of a common state space \mathcal{S} .

⁶¹ For simplicity, we assume that the measure M over $\Delta(\mathcal{S})$ has finite support.

The individual evaluates this lottery again using the function $V(\cdot)$, which yields the preference function

$$W(f(\cdot)) \equiv V(\dots; CE(f, \mu_k), M(\mu_k); \dots) \quad (13.7)$$

Together, this implies the chain of indifference

$$\begin{aligned} f(\cdot) &\sim (\dots; f(\cdot), \mu_k; \dots) \sim (\dots; CE(f(\cdot), \mu_k), \mu_k; \dots) \sim \\ &(\dots; CE(f(\cdot), \mu_k), M(\mu_k); \dots) \end{aligned}$$

or equivalently

$$\begin{aligned} (\dots; x_j, E_j; \dots) &\sim (\dots; (\dots; x_j, E_j; \dots), \mu_k; \dots) \\ &\sim (\dots; CE((\dots; x_j, E_j; \dots), \mu_k), \mu_k; \dots) \\ &\sim (\dots; CE((\dots; x_j, E_j; \dots), \mu_k), M(\mu_k); \dots) \end{aligned}$$

where the first object is the subjective act $f(\cdot)$ from states to payoffs which is being evaluated, the second is the representation of this act as a map from the set of measures $\mu_k \in \Delta(\mathcal{S})$ to subjective acts which takes the value $f(\cdot)$ for all μ_k , the third is the map from this set of measures to the act's certainty equivalent under each measure, and the fourth is the objective lottery over these certainty equivalents values implied by the probability measure $M(\cdot)$ over the measures.

If the preference function $V(\cdot)$ over objective lotteries is expected utility, then the final evaluation reduces to subjective expected utility. However, Segal suggests using *anticipated utility*, or *rank-dependent expected utility* form (Quiggin (1982), Yaari (1987), Segal (1987b)). For a lottery $\mathbf{P} = (\dots; \alpha_i, p_i; \dots)$, with $\alpha_1 > \dots > \alpha_n$, this preference function takes the form

$$V(\mathbf{P}) = U(\alpha_1) \cdot \delta(p_1) + \sum_{i=2}^n U(\alpha_i) \cdot [\delta(p_1 + \dots + p_i) - \delta(p_1 + \dots + p_{i-1})].$$

As such, it is analogous to the Choquet expected utility form (13.5) over subjective acts, with the nonadditive capacity ν replaced by a nonlinear (and typically convex) *distortion function* $\delta: [0,1] \rightarrow [0,1]$ for which $\delta(0) = 0$ and $\delta(1) = 1$. When $\delta(\cdot)$ is linear, this form reduces to objective expected utility.

Labeling the induced lotteries $\mathbf{P}_k = (\dots; x_j, \mu_k(E_j); \dots)$ so that $V(\mathbf{P}_1) \geq \dots \geq V(\mathbf{P}_n)$, and observing that the certainty equivalent of a lottery \mathbf{P} is given by $U^{-1}(V(\mathbf{P}))$, the Recursive preference function over subjective acts takes the form

$$\begin{aligned} W(f(\cdot)) &= V(\mathbf{P}_1) \cdot \delta(M(\mu_1)) + \sum_{i=2}^n V(\mathbf{P}_i) \cdot [\delta(M(\{\mu_1, \dots, \mu_i\})) \\ &\quad - \delta(M(\{\mu_1, \dots, \mu_{i-1}\}))]. \end{aligned}$$

To illustrate how this model can accommodate Ellsberg preferences in the Three-Color Urn, keep the normalization $U(\$100) = 1$ and $U(\$0) = 0$, and suppose that the decision-maker assigns probability $1/2$ to each of the distributions

$$\mu_1(s_r, s_b, s_y) = (1/3, 2/3, 0) \quad \mu_2(s_r, s_b, s_y) = (1/3, 0, 2/3). \quad (13.8)$$

Then bet a_1 on red induces the same lottery $\mathbf{P} = \{\$100; 1/3; \$0, 2/3\}$ under either measure, so we can identify it with the degenerate two-stage lottery $\{\mathbf{P}, 1\}$ to obtain $W(a_1) = V(\mathbf{P}) = \delta(1/3)$. A similar argument establishes $W(a_4) = V(\mathbf{P}) = \delta(2/3)$.

Consider now the bet a_2 on black. Under distribution μ_1 it yields the lottery $\mathbf{P}_1 = \{\$100, 2/3, \$0, 1/3\}$, while under μ_2 it yields $\mathbf{P}_2 = \{\$100, 0; \$0, 1\}$. We have $V(\mathbf{P}_1) = \delta(2/3)$ and $V(\mathbf{P}_2) = 0$, so from equation (13.7), $W(a_2) = V(\mathbf{P}_1) \cdot \delta(M(\mu_1)) = \delta(2/3) \cdot \delta(1/2)$. A similar argument establishes $W(a_3) = 1 \cdot \delta(1/2) + \delta(1/3)[1 - \delta(1/3)]$. For distortion function $\delta(p) = (e^p - 1)/(e - 1)$ these values become⁶² $W(a_1) = 0.2302 > 0.2082 = W(a_2)$ and $W(a_3) = 0.5208 < 0.5516 = W(a_4)$, which yield the Ellsberg rankings.

While Segal (1987a) utilizes the same preference function to evaluate single-stage lotteries as well as lotteries over certainty equivalents, Segal (1990) considers the possibility of using different preference functions in the two stages.

13.5.4 Klibanoff, Marinacci and Mukerji's Smooth Ambiguity Preferences Model

The *Smooth Ambiguity Preferences Model* Klibanoff et al. (2005) (*KMM*) follows a similar two-stage approach. The main differences with Segal's Recursive Model are that the primitives are preferences over horse lotteries, i.e., subjective acts, rather than objective roulette lotteries, and that different preference functions are used in the different stages. In particular, these are each expected utility functions, with different utility functions. This is related to Example 2 in Segal (1990).

In applications, the smooth representation typically takes the following form: for any act (horse lottery) $f : \mathcal{S} \rightarrow \mathcal{X}$, the smooth ambiguity index is

$$W(f(\cdot)) = \int_{\Delta(\mathcal{S})} \phi \left(\int U(f(\cdot)) d\mu \right) dM(\mu) \quad (13.9)$$

where $U(\cdot)$ is the von Neumann-Morgenstern utility function, M is the individual's *second-order prior*, that is, a probability measure over the set $\Delta(\mathcal{S})$ of measures μ on the state

⁶² Segal (1987a, p.185), see also his Theorem 4.2 for a general result along the lines of this example.

space \mathcal{S} , and $\varphi(\cdot)$ is the individual's *second-order utility function*. As in Segal's Recursive model, the interpretation is that the individual is uncertain about which probability measure μ describes the likelihood of realizations of the state; this second-order uncertainty is modeled using the second-order prior M . Each measure μ determines a different expected utility $\int U(f(\cdot))d\mu$ or the act f . If the individual simply averaged these expected utilities using M , and hence evaluated f according to

$$\int_{\Delta(\mathcal{S})} \int_{\mathcal{S}} U(f(\cdot)) d\mu dM(\mu)$$

this would yield expected utility preferences, where the subjective probability measure is the "average measure" that assigns to each event E the probability $\int \mu(E)dM(\mu)$. Instead, the individual "distorts" the expected utilities $\int U(f(\cdot))d\mu$ using the second-order utility function $\varphi(\cdot)$. As we shall see momentarily, concavity of $\varphi(\cdot)$ corresponds to a notion of ambiguity aversion. (Similarly, convexity of $\varphi(\cdot)$ corresponds to ambiguity preference, and linearity of $\varphi(\cdot)$ corresponds to expected utility and hence ambiguity neutrality.)

To see how this representation can accommodate the modal preferences in the Three-Color Paradox, suppose that, as in the preceding subsection, the individual assigns equal second-order probabilities to the distributions in equation (13.8). Continue to assume that $U(\$100) = 1$ and $U(\$0) = 0$, and let $\varphi(r) = \sqrt{r}$. Then $W(a_1) = \varphi(1/3) = 0.5773 > 0.4082 = 1/3 \cdot \varphi(1/3) + 1/2 \cdot \varphi(0) = W(a_2)$ and similarly $W(a_3) = 1/2 \cdot \varphi(1/3) + 1/2 \cdot \varphi(1) = 0.7887 < 0.8165 = \varphi(2/3) = W(a_4)$, as required.

We now show how *KMM* obtain the representation in equation (13.9); in so doing, we also illustrate the formal connection with Segal's approach.

KMM adopt a decision framework in which, loosely speaking, the realizations of objective roulette wheels are modeled as part of the state space; preferences over acts that depend *only* on the roulette-wheel component of the state space are assumed consistent with *EU* with a uniform prior probability measure. Formally, the state space \mathcal{S} is the Cartesian product of a set Ω of interest and the interval $[0,1]$ (the roulette wheel). The set of prizes is an interval of real numbers. The individual is endowed with not one, but two preference relations. The first, denoted \succsim , is a preference over (standard) acts $f : \mathcal{S} \rightarrow \mathcal{X}$. Its restriction to the set of acts that only depend on the roulette-wheel component $[0,1]$ of the state space \mathcal{S} is assumed to admit an expected utility representation, with a continuous, strictly increasing utility U and a prior λ which is the Lebesgue measure on $[0,1]$. It is clear that, for every roulette lottery $\mathbf{P} = (\dots; x_i, p_i; \dots)$ with prizes in the set \mathcal{X} , there exists an act $f_{\mathbf{P}}$ that depends only on the roulette-wheel component of the state space and which generates the same probability distribution as \mathbf{P} , in the sense that $\lambda(\{r : f_{\mathbf{P}}(\omega, r) = x_i \text{ for all } \omega\}) = p_i$ for all i . For this reason, call such acts *lottery acts*.

The second preference relation, denoted $\succsim^{(2)}$, is defined over *second-order acts*, i.e., maps $f^{(2)} : \Delta(\mathcal{S}) \rightarrow \mathcal{X}$. We return to the interpretation of this relation below. The relation $\succsim^{(2)}$ is assumed to have an expected utility representation, with continuous, strictly increasing utility v (not necessarily equal to U) and probability measure $M \in \Delta(\Delta(\mathcal{S}))$.

The key axiom, *Consistency*, connects the preferences \succsim on standard acts and $\succsim^{(2)}$ on second-order acts. It also indicates how the Smooth Ambiguity Preferences model is related to Segal’s Recursive approach. The Consistency axiom formalizes Segal’s analysis of the Three-Color Urn in the previous subsection (and is also consistent with Segal’s own intended use of the model). We assumed that each act a_1, a_2, a_3, a_4 was associated with a collection of single-stage lotteries, corresponding to the probability distributions over states that the individual contemplates; then, we evaluated the act by evaluating the two-stage lottery obtained by attaching probabilities to each such single-stage lottery. *KMM’s* Consistency axiom makes this process explicit: they associate with each standard act f the second-order act $f^{(2)}$ such that, for every $\mu \in \Delta(\mathcal{S})$, $f^{(2)}(\mu)$ is a certainty equivalent of the lottery act that yields each outcome x in the range of f with probability measure $\mu(f^{-1}(x))$. Given the assumption that preferences over lottery acts are consistent with expected utility, with a continuous and strictly increasing (hence, invertible) utility function $U(\cdot)$, this means that

$$f^{(2)}(\mu) = U^{-1} \left(\int U(f(\cdot)) d\mu \right)$$

for every $\mu \in \Delta(\mathcal{S})$. The key assumption is then

Consistency: For all acts f, g , $f \succsim g$ iff $f^{(2)} \succsim^{(2)} g^{(2)}$.

That is: the ranking of f and g is precisely the ranking of the corresponding second-order acts $f^{(2)}$ and $g^{(2)}$.

Consistency and the assumption that preferences over second-order acts admit an expected utility representation then immediately imply that standard, first-order acts f can be ranked according to

$$W(f(\cdot)) = \int_{\Delta(\mathcal{S})} v(f^{(2)}(\mu)) dM(\mu) = \int_{\Delta(\mathcal{S})} \phi \left(\int U(f(\cdot)) d\mu \right) dM(\mu)$$

where $\varphi(\cdot) = v(U^{-1}(\cdot))$ is the *second-order utility* and M the *second-order probability measure*, or *second-order prior*. This is the representation in equation (13.9). (In *KMM*, each measure μ in the support of the second-order prior is an independent product of Lebesgue measure λ on $[0,1]$ and a suitable probability measure on the coordinate of actual interest, namely Ω . In applications, the Lebesgue component is simply omitted.) Notice that, while this model is commonly referred to as the *Smooth Ambiguity Preferences* model, the function φ is not necessarily smooth (though it can be, and typically is, in applications).

The interpretation of second-order acts requires some elaboration. *KMM* acknowledge that a direct interpretation of them as bets on different elements in $\Delta(\mathcal{S})$ is not straightforward:

“there is a question whether preferences with respect to these acts are observable. The mapping from observable events to events in $\Delta(\mathcal{S})$ may not always be evident. When it is not evident we may need something richer than behavioral data, perhaps cognitive data or thought experiments” (p. 1854).

If second-order acts are not directly observable, then the second-order utility v and second-order probability measure M should be thought of solely as a way to provide a numerical representation of preferences over *first-order*, standard acts. The question is then whether these objects can be identified on the basis of preferences over standard acts alone. In general, the answer is negative: for a trivial example, observe that if the second-order utility $\phi(\cdot)$ is linear (equivalently, if $v = U$), one can equivalently replace a given second-order probability measure M with one that is concentrated on the “average probability measure” $\int \mu dM(\mu)$ without changing first-order preferences at all. As *KMM* indicate, it is an open question to what extent and under what conditions $\phi(\cdot)$ and M can be identified from first-order preferences alone. For an environment in which uncertainty concerns infinitely many repetitions of an experiment, a positive answer is provided by Klibanoff et al. (2011) under a “symmetry” restriction on preferences. A related result can be found in Al-Najjar and De Castro (2013).

On the other hand, there are settings in which second-order acts are observable. An obvious example are experiments involving physical devices: one can resolve bets on the composition of an urn by simply opening the urn and examining its contents. Bets on long-run averages can also be thought of as observable second-order acts, though arguably only as idealizations or approximations. However, in these cases, first- and second-order acts only differ because they depend on different “coordinates” of the state space; the model in equation (13.9) essentially allows the individual to be differently risk-averse, depending on the *source* of uncertainty. In other words, rather than ambiguity attitudes, the smooth model reflects *issue preference*, which we consider in the following section.

An alternative axiomatization of the Smooth Ambiguity Preferences model is provided by Seo (2009). That paper takes a different approach. The individual is assumed to have preferences over roulette-horse-roulette acts (see Section 13.3.4). The key assumption is that individuals do *not* reduce compound lotteries, though they satisfy a suitable “dominance” property. Seo obtains an extension of the representation in equation (13.9) to lotteries over horse-roulette acts; the function $\phi(\cdot)$ captures both attitudes towards ambiguity and violations of reduction. Thus, one implication of his model is that an agent who reduces compound *objective* lotteries must be ambiguity-neutral. There is some evidence that attitudes toward reduction and ambiguity are correlated (Halevy (2007)); however, one may wish for a model that does not impose perfect correlation between these distinct behavioral traits.

The main attraction of the Smooth Ambiguity Preferences model is its tractability. The numerical example above suggests that calculations may involve simple adaptations of familiar expected utility analysis. As further evidence, *KMM* show that concavity of the second-order utility function $\phi(\cdot)$ corresponds to a notion of “ambiguity aversion” which implies Schmeidler’s Uncertainty Aversion Axiom (although their notion is defined in terms of second-order beliefs, and hence subject to the preceding qualifications).

13.5.5 Ergin and Gul’s Issue-Preference Model

Ergin and Gul (2009) propose a variant of Segal’s two-stage approach, and provide an elegant axiomatization in a Savage-style environment. Their model highlights a different interpretation of the Ellsberg Paradoxes.

To illustrate, it is useful to rephrase the Three-Color Paradox slightly. Imagine that the balls in the urn are numbered from 1 to 90. Balls numbered 1–30 are known to be red. Balls numbered 31–90 may be either black or yellow: the color of each of these balls is not known. Thus, the experiment involves two *sources of uncertainty*, or *issues*: the number of the ball drawn, and the colors of balls 31–90. Importantly, these issues are *independent*: if you learn the number of the ball drawn, that does not tell you anything about the colors of balls 31–90, or conversely.

This richer description can be modeled using a two-coordinate state space, for example

$$\mathcal{S} = \{1, \dots, 90\} \times \{c = (c_{31}, \dots, c_{90}) : \forall i, c_i \in \{b, \gamma\}\}.$$

Denote an arbitrary state in \mathcal{S} by $s = (n, c)$, where n is the number of the drawn ball and c is a color assignment, i.e., the vector (c_{31}, \dots, c_{90}) . The bets in the Three-Color Paradox can then be modeled as follows (with some abuse of notation, we use the same letters as in Section 13.1): the bets on “red” and “black or yellow” correspond respectively to

$$a_1((n, c)) = \begin{cases} \$100 & n = 1, \dots, 30 \\ \$0 & n \geq 31 \end{cases} \quad \text{and} \quad a_4((n, c)) = \begin{cases} \$0 & n = 1, \dots, 30 \\ \$100 & n \geq 31 \end{cases}.$$

The bets on “black” and “red or yellow,” on the other hand, correspond to

$$a_2((n, c)) = \begin{cases} \$0 & n = 1, \dots, 30 \\ \$100 & n \geq 31 \text{ and } c_n = b \\ \$0 & n \geq 31 \text{ and } c_n = \gamma \end{cases} \quad \text{and} \quad a_3((n, c)) = \begin{cases} \$100 & n = 1, \dots, 30 \\ \$0 & n \geq 31 \text{ and } c_n = b \\ \$100 & n \geq 31 \text{ and } c_n = \gamma. \end{cases}$$

An obvious difference between a_1 and a_4 on one hand, and a_2 and a_3 on the other, is that the former only depend upon the number of the ball drawn, whereas the latter depend upon both the number of the ball drawn and the color assignment to balls 31...90 (in particular, the color of the ball drawn). Now suppose that the individual has *probabilistic beliefs*, and in particular believes that every ball is equally likely to be drawn, and also that, independently, every color assignment is equally likely to have occurred. This implies that, for both acts a_1 and a_2 , the probability of receiving \$100 is $1/3$; similarly, the probability of winning is $2/3$ for both acts a_3 and a_4 . The Ellsberg preferences $a_1 \succ a_2$ and $a_3 \prec a_4$ can then be interpreted as follows: despite the fact that the winning probabilities in each comparison are the same, the individual prefers bets that only depend upon the first issue (here, the number of the ball drawn) to bets that depend upon both issues (the number as well as the color assignment).

Recall that an individual is probabilistically sophisticated if he or she evaluates an act $f = (\dots; x_j, E_j; \dots)$ by first reducing it to a lottery $\mathbf{P}_f = (\dots; x_j, \mu(E_j); \dots)$ where μ is their subjective probability over the state space \mathcal{S} , and then applying a suitable preference function over lotteries. Even in the above modified formulation of the Three-Color Urn bets, Ellsberg preferences are not probabilistically sophisticated—acts that induce the same probability distribution over prizes are strictly ranked. However, Ergin and Gul note that they are *second-order probabilistically sophisticated*, in the sense that each act is evaluated by first reducing it to a *compound* lottery, in which the first stage corresponds to the assignment of colors to balls 31, ..., 90 and the second stage corresponds to the draw of a numbered ball from the urn, and then applying a suitable *compound-lottery* preference function. This reduction of acts to compound lottery is essentially what was done in order to apply Segal's two-stage model to the Three-Color Paradox. Ergin and Gul point out that this form of reduction is a higher-order version of probabilistic sophistication and, more importantly, provide a complete axiomatization of the resulting preferences.

Their axiomatization builds upon that of Machina and Schmeidler (1992) for probabilistically sophisticated preferences in a Savage framework (Section 3). To provide some detail, assume that the state space is a Cartesian product, $\mathcal{S} = \mathcal{S}_a \times \mathcal{S}_b$, where the two coordinates correspond to *issue a* and *issue b*. In the Three-Color Urn, issue a is the number of the ball drawn, and issue b is the color assignment. The Machina-Schmeidler Strong Comparative Probability axiom is applied only to acts that depend solely on issue a . Ergin and Gul's key novel axiom, *a|b Strong Comparative Probability*, serves three purposes. First, it implies that Strong Comparative Probability also holds for acts that depend solely on issue b . Second, it implies that the two issues are statistically independent, as in the example. Third, it implies, loosely speaking, that the individual is concerned with second-order risk related to issue b , and not second-order risk related to issue a (Ergin and Gul (2009, p.906) provide an illustration).

To formally describe the Ergin–Gul representation, note that exploiting the product structure of the state space allows an act f to be represented as

$$f = \begin{pmatrix} x_{11} & \cdots & x_{1m} & A_1 \\ \vdots & \ddots & \vdots & \vdots \\ x_{n1} & \cdots & x_{nm} & A_n \\ B_1 & \cdots & B_m & \end{pmatrix}$$

where $\{A_1, \dots, A_n\}$ is a partition of \mathcal{S}_a , $\{B_1, \dots, B_m\}$ is a partition of \mathcal{S}_b , and x_{ij} is the prize delivered by f at all states $s \in A_i \times B_j$. Given a product measure $\mu = \mu_a \times \mu_b$ on \mathcal{S} , associate with f the two-stage compound lottery

$$\mathbf{P}_f^* = (\dots; (\dots; x_{ij}, \mu_a(A_i)), \mu_b(B_j); \dots).$$

Under Ergin and Gul's axioms, there exists a unique product measure $\mu = \mu_a \times \mu_b$ on \mathcal{S} such that every act f is evaluated according to $W(f(\cdot)) = V^*(\mathbf{P}_f^*)$, where \mathbf{P}_f^* is the compound lottery associated with f and V^* is a suitably continuous preference function on two-stage lotteries that satisfies stochastic dominance; again, see the original paper for details. If a version of Savage's Sure-Thing Principle is imposed, this function reduces to two-stage expected utility with different utility functions, as in the Smooth Ambiguity Preferences model of Section 13.5.3. If instead a *Comonotonic Sure-Thing Principle* axiom is imposed, the first-stage lottery is evaluated using a Choquet integral. Ergin and Gul also draw a formal connection between uncertainty aversion and issue preference.

Nau (2006) proposes a related model, which allows for state-dependent preferences but restricts attention to additive representations in both stages. Different perspectives on source preference are suggested by Chew and Sagi (2008) and Abdellaoui et al. (2011). Finally, Skiadas (2013) assumes that prizes are monetary and preferences over acts are scale-invariant: $(\dots; x_j, E_j; \dots) \succ (\dots; y_j, F_j; \dots)$ if and only if, for every real number $\alpha > 0$, $(\dots; \alpha \cdot x_j, E_j; \dots) \succ (\dots; \alpha \cdot y_j, F_j; \dots)$. This yields a two-stage expected utility representation *a la* Ergin and Gul, with different constant relative risk averse utility functions for the two stages.

13.5.6 Vector Expected Utility

Siniscalchi (2009a) proposes a different approach to modeling ambiguity-sensitive preferences. There are three key ingredients in this approach. First, the individual evaluates acts by a process reminiscent of “anchoring and adjustment” (Tversky and Kahneman (1974)). Second, the *anchor* is a standard expected utility evaluation, where expectations

are taken with respect to a *baseline* subjective probability measure. Third, the *adjustment* depends upon a measure of statistical similarity of the act under consideration with “examples” or “models” of ambiguous acts.

The proposed representation is called *vector expected utility (VEU)* because a key role is played by a vector of expectations. Specifically, the *VEU* representation evaluates acts according to

$$W(f(\cdot)) = \int U(f(\cdot)) \cdot d\mu + A(\text{cov}(U(f(\cdot)), \zeta_1), \dots, \text{cov}(U(f(\cdot)), \zeta_n)) \tag{13.10}$$

where μ is the baseline prior, ζ_1, \dots, ζ_n are *adjustment factors*, i.e., real-valued functions on the state space \mathcal{S} that satisfy $\int \zeta_i d\mu = 0$, $\text{cov}(\cdot, \cdot)$ denotes covariance (computed with respect to μ), and $A(\cdot)$ is the *adjustment function*, which satisfies $A(0, \dots, 0) = 0$ and $A(\varphi) = A(-\varphi)$ for every vector φ . Notice that, due to the normalization of the adjustment factors ζ_1, \dots, ζ_n ,

$$\text{cov}(U(f(\cdot)), \zeta_i) = \int U(f(\cdot)) \cdot \zeta_i \cdot d\mu.$$

To see how the *VEU* model accommodates Ellsberg preferences, consider a simple specification with a single adjustment factor ζ_1 , defined by

$$\zeta_1(s_r) = 0 \quad \zeta_1(s_b) = 1 \quad \zeta_1(s_y) = -1$$

a uniform baseline prior, and the adjustment function

$$A(\varphi) = -|\varphi|.$$

The adjustment factor ζ_1 can be interpreted as a bet whereby the individual wins or loses one util, depending on which of the two ambiguous colors is drawn, but nothing happens if the unambiguous color red is drawn.⁶³ Intuitively, this suggests that ζ_1 is a “purely ambiguous” bet, and indeed a model of the kind of bets that are ambiguous in this problem.⁶⁴

A related interpretation of the adjustment factor ζ_1 formalizes the following intuition, proposed by Epstein and Zhang (2001) for the modal preferences in the Ellsberg Paradox:

*“The intuition for this reversal is the **complementarity** between s_b and s_y —there is imprecision regarding the likelihood of s_b , whereas $\{s_b, s_g\}$ has precise probability 2/3” (p. 271; **boldface** added for emphasis, and notation adapted to this chapter).*

⁶³ For the sake of this interpretation, assume that initial utility is constant and normalized to zero.

⁶⁴ This intuitive interpretation corresponds closely to the way adjustment factors are formally identified from preferences: see Section 4.1 in Siniscalchi (2009a).

The fact that $\zeta_1(s_b) = 1$ whereas $\zeta_1(s_r) = 0$ reflects the perception that s_b is ambiguous but s_r is not; setting $\zeta_1(s_y) = -1$ indicates that ambiguity about s_b and s_y “cancels out.” In this sense, s_b and s_y are *complementary* in the *VEU* representation; this formalizes the Epstein-Zhang interpretation. (The relationship with the previous interpretation is that, in order for ζ_1 to be a purely ambiguous bet in this example, it *must* reflect the complementarity between s_b and s_y .)

We now verify that the *VEU* representation proposed here induces the Ellsberg rankings.

$$\begin{aligned} W(a_1) &= 1/3 - |0| = 1/3 \\ W(a_2) &= 1/3 - |1/3| = 0 \\ W(a_3) &= 2/3 - |-1/3| = 1/3 \\ W(a_4) &= 2/3 - |0| = 2/3 \end{aligned}$$

In the calculation of $W(a_4)$, the covariance term is zero precisely because ambiguity about s_b and s_y cancels out: $\text{cov}(U(a_4), \zeta_1) = 1/3 \cdot \sum_s \zeta_1(s) a_4(s) = 1/3 \cdot (0 \cdot 0 + 1 \cdot 1 + 1 \cdot (-1)) = 0$.

The key axiom in the characterization of *VEU* preferences is *Complementary Independence*. A preliminary definition: two acts f, f^c are *complementary* if the mixture $1/2 \cdot f + 1/2 \cdot f^c$ is a constant act. Intuitively, complementary acts are the “negative” of each other, up to a constant. An important feature is that they share the same *prize variability*, or volatility: if the (utility of the) prizes delivered by f at different states varies a lot, the same is true of prizes delivered by f^c , except that prizes vary in opposite directions. The main axiom for *VEU* preferences is then

Complementary Independence: For all acts f, g, f^c, g^c such that f, f^c and respectively g, g^c are complementary: if $f \succcurlyeq f^c$ and $g \succcurlyeq g^c$, then $1/2 \cdot f + 1/2 \cdot g \succcurlyeq 1/2 \cdot f^c + 1/2 \cdot g^c$.

When comparing complementary acts, variability cannot play a role because, as just noted, any two complementary acts have the same volatility. The axiom requires that, absent volatility considerations, a form of Independence must hold. This implies that any departures from expected utility must be determined by *attitudes toward variability*. For further interpretation, the reader is referred to the Siniscalchi (2009a).

VEU preferences can accommodate a wide range of ambiguity attitudes. If the Uncertainty Aversion Axiom is imposed, the function A must be negative and concave. It turns out that a negative, but not necessarily concave function A also captures a notion of ambiguity aversion; we discuss this further in Section 13.6.

Chambers et al. (2013) consider preferences that admit a representation similar to that of equation (13.10), but with an arbitrary aggregator to combine the baseline

expectation and adjustment terms. They adopt the Complementary Independence axiom, and also require a form of uncertainty aversion. See also Grant and Polak (2013).

13.5.7 Variational and Multiplier Preferences

The Uncertainty Aversion Axiom of Section 13.5.1 is common to a broad class of preference models. In addition to *MEU* preferences, suitable parameterizations of smooth ambiguity, Ergin–Gul, and *VEU* preferences also satisfy it. As we will discuss in Section 13.6, its interpretation is not entirely straightforward, or uncontroversial; on the other hand, the fact that it is a familiar and analytically convenient quasiconcavity assumption has made it an obvious reference point in the theoretical and applied literature on ambiguity-sensitive preferences. This subsection and the next two discuss other preference models that satisfy Uncertainty Aversion, and can be thought of as direct generalizations of the *MEU* specification.

To introduce these models, recall that the characterization of *MEU* preferences also involves Certainty Independence. Unlike the Uncertainty Aversion Axiom, Certainty Independence is seen to be quite restrictive. We noted in Section 13.5.1 that, when an act such as a_2 (bet on black) in the Three-Color-Urn is mixed with a constant act, no hedging occurs: like the act a_2 , the mixture yields a strictly better prize in case a black ball is drawn than if a yellow or red ball is drawn. By way of contrast, mixing a_2 with a bet on yellow results in an act which delivers a good prize if either of the ambiguous colors is drawn. However, mixing a_2 with a constant act does have two potentially important effects. First, it reduces the *utility variability* across the two ambiguous states s_b and s_y : for instance, a_2 yields utility $U(\$100)$ and $U(\$0)$ in case a black or, respectively, nonblack ball is drawn, whereas a 50:50 mixture of a_2 with the constant act $\$0$ yields $1/2 \cdot U(\$100) + 1/2 \cdot U(\$0)$ and $U(\$0)$ respectively. Second, mixing with a constant act also shifts utilities up or down: to see this, compare the utilities delivered by a 50:50 mixture of a_2 with $\$0$ to those delivered by a 50:50 mixture of a_2 with $\$100$, namely $U(\$100)$ for a black ball and $1/2 \cdot U(\$100) + 1/2 \cdot U(\$0)$ for a nonblack ball.

Both of these changes may potentially affect preferences. After all, in choice under risk, changes in the *scale* (i.e., outcome variability) and *location* (upward/downward shifts) of *monetary prizes* can affect preferences. For instance, if preferences exhibit decreasing absolute risk aversion, an individual becomes more willing to take bets as their initial (nonrandom) wealth increases, i.e., shifts upward. Similarly, unless preferences display constant relative risk aversion, an individual with constant initial wealth may accept a given bet, but turn it down if both their wealth and the stakes in the bet are doubled (or vice versa). Arguably, it would be desirable to allow for similar behavior in the context of ambiguity: changing the scale and/or location of utilities may influence preferences.

However, Certainty Independence is a location- and scale-invariance assumption, so such phenomena are ruled out.

13.5.7.1 Variational Preferences

Motivated in part by these considerations, Maccheroni et al. (2006a) propose a weakening of Certainty Independence:

Weak Certainty Independence: For all acts f, g , constant acts x, y and all $\alpha \in (0, 1]$: if $\alpha \cdot f + (1 - \alpha) \cdot x \succcurlyeq \alpha \cdot g + (1 - \alpha) \cdot x$, then $\alpha \cdot f + (1 - \alpha) \cdot y \succcurlyeq \alpha \cdot g + (1 - \alpha) \cdot y$.

This axiom assumes invariance to changes in the location, but *not* the scale, of utilities. In the Three-Color Urn problem, it allows for the following pattern of preferences:

$$a_1 > a_2 \quad \text{and} \quad 1/2 \cdot a_1 + 1/2 \cdot 0 < 1/2 \cdot a_2 + 1/2 \cdot 0.$$

A possible interpretation might be that, by scaling down the utilities of the acts involved in the comparison, the individual becomes less sensitive to the fact that black is an ambiguous color but red is not. (We do not claim that such preferences would be reasonable or natural, merely that they are allowed by Weak Certainty Independence.) At the same time, Maccheroni et al.'s axiom still requires that, since the 50:50 mixture of a_1 with \$0 is considered worse than the 50:50 mixture of a_2 with \$0, the same must be true if \$0 is replaced with \$100.

Maccheroni et al. show that the Weak Order, Non-Degeneracy, Monotonicity, Continuity, Uncertainty Aversion and Weak Certainty Independence Axioms are necessary and sufficient for the existence of a representation of the form

$$W(f(\cdot)) \equiv \min_{\mu \in \Delta(\mathcal{S})} \int U(f(\cdot)) d\mu + c(\mu) \quad (13.11)$$

where $c : \Delta(\mathcal{S}) \rightarrow [0, \infty]$ is a lower semicontinuous, convex *cost function* that satisfies $\inf_{\mu} c(\mu) = 0$. They call equation (13.11) the *variational representation* of preferences.

The cost function $c(\cdot)$ is the key ingredient of the representation. Once a utility function has been fixed, a “canonical” or “minimal” cost function is uniquely identified (refer to Maccheroni et al. for details). If the utility function U is replaced with a positive affine transformation $\alpha \cdot U + \beta$, the cost function c must be replaced with $\alpha \cdot c$.

Observe that the minimization in equation (13.11) is over all priors over the state space \mathcal{S} . However, one can obtain *MEU* preferences as a special case as follows: fix a compact, convex set C of priors, and define

$$c(\mu) = \begin{cases} 0 & \mu \in C \\ \infty & \mu \notin C \end{cases}$$

(such a function is often called an “indicator function” in convex analysis). With this specification of the cost function, it is clear that a prior $\mu \notin C$ will never be a solution

of the minimization problem in equation (13.11) whereas priors in \mathcal{C} are costless, so the minimization problem reduces to that in equation (13.1).

Maccheroni et al. propose the following “malevolent Nature” interpretation. When the individual contemplates choosing an act f , Nature tries to minimize its expected utility. To do so, she can choose any prior $\mu \in \Delta(\mathcal{S})$; however, to choose μ , Nature must pay a cost $c(\mu)$. In the case of *MEU* preferences, priors in \mathcal{C} are costless, and all other priors are infinitely costly; however, the variational representation allows for intermediate cases as well—that is, for priors that have a nonzero but finite cost. Nature’s problem is thus to minimize the sum of the individual’s expected utility and the cost of priors.

Skiadas (2013) also axiomatizes a version of Variational Preferences with constant relative risk averse utility.

13.5.7.2 Multiplier Preferences

Another special parameterization of Variational Preferences deserves special mention. Hansen and Sargent (2001) and Hansen et al. (1999) consider the following specification:

$$W(f(\cdot)) \equiv \min_{\mu \in \Delta(\mathcal{S})} \int U(f(\cdot)) d\mu + \theta \cdot R(\mu || \mu^*) \tag{13.12}$$

where attention is restricted to countably additive probability measures when the state space is infinite, $\theta \geq 0$ is a parameter, and $R(\mu || \mu^*)$ denotes the *relative entropy* of the prior μ with respect to a *reference prior* μ^* . The latter is defined as

$$R(\mu || \mu^*) = \int \left(\log \frac{d\mu}{d\mu^*} \right) d\mu$$

if μ is absolutely continuous with respect to μ^* , and $R(\mu^* || \mu^*) = \infty$ otherwise. This defines the so-called *Multiplier Preferences* preference function. Equation (13.12) is a special case of equation (13.11), yielding tractability since the explicit solution to the problem in equation (13.12) is known, and indeed one can show that

$$W(f(\cdot)) = -\theta \cdot \log \left\{ - \int - \exp \left(- \frac{U(f(\cdot))}{\theta} \right) d\mu^* \right\}$$

(the reason for the two consecutive minus signs will be clear momentarily). An exact characterization of Multiplier Preferences within the class of Variational Preferences is provided by Strzalecki (2011).

The above equation indicates that, if one restricts attention to horse lotteries, i.e. acts that map states to prizes (rather than to roulette lotteries over prizes), Multiplier Preferences are ordinally equivalent to *SEU* preferences with utility function $-\exp(-\theta^{-1} \cdot U(\cdot))$. To see this, note that $W(f(\cdot))$ is the composition of the function $f - \exp(-\theta^{-1} \cdot U(\cdot)) d\mu$

which takes on strictly negative values, with the function $x \rightarrow -\theta \log(-x)$, which is strictly increasing when restricted to the negative reals. On the other hand, if one restricts attention to roulette lotteries, the individual's risk preferences are represented by $U(\cdot)$. In other words, one can interpret the multiplier-preferences model as postulating that the individual is more risk-averse toward *subjective* risk than toward *objective* risk. The idea that risk attitudes may be *source-dependent* was discussed in Section 13.5.5.

13.5.8 Confidence-Function Preferences

Chateauneuf and Faro (2009) propose a model that is “dual” to Variational Preferences: whereas the latter relax the scale-invariance requirement of Certainty Independence but preserve location-invariance, Chateauneuf and Faro's *Confidence-Function Preferences* retain scale invariance but drop translation invariance. More precisely, these authors assume that the set \mathcal{X} of prizes contains a worst element x^* , and impose

Worst Independence: For all acts f, g and all $\alpha \in (0, 1]$: $f \sim g$ implies $\alpha \cdot f + (1 - \alpha) \cdot x^* \sim \alpha \cdot g + (1 - \alpha) \cdot x^*$.

Together with Weak Order, Continuity, Monotonicity, Ambiguity Aversion and other axioms, Worst Independence characterizes the following representation:

$$W(f(\cdot)) = \min_{\mu: \varphi(\mu) \geq \alpha} \frac{1}{\varphi(\mu)} \int U(f(\cdot)) d\mu \quad (13.13)$$

where the utility function $U(\cdot)$ satisfies $U(x^*) = 0$, α is a parameter in $(0, 1]$, and $\varphi(\cdot)$ is an upper semicontinuous, quasiconcave function on $\Delta(\mathcal{S})$ that satisfies $\varphi(\mu) = 1$ for some prior μ .

The proposed interpretation is that $\varphi(\cdot)$ measures the “degree of confidence” that the individual assigns to the different measures. Nature is malevolent, and wants to minimize the individual's expected utility. However, she must choose only priors to which the individual assigns a high enough level of confidence. Furthermore, Nature must pay a multiplicative cost $\varphi(\mu)^{-1}$ to choose the prior μ .

Confidence-function preferences are positively homogeneous: if all utility values are multiplied by some constant $\beta > 0$, the function in equation (13.13) is also multiplied by β . If utility is also positively homogeneous, i.e., if it displays constant relative risk aversion, then the overall representation of preferences is also positively homogeneous. This is convenient in certain applications, especially in financial economics. Furthermore, it makes it possible to invoke standard results on the existence of a representative agent.

13.5.9 Uncertainty-Averse Preferences

Cerreia-Vioglio et al. (2011a) characterize the class of preferences that satisfy Continuity, Monotonicity and Uncertainty Aversion, in addition to Independence on lotteries and

certain other technical assumptions. In other words, they drop both location and scale invariance. Their powerful result unifies all the models that generalize *MEU*, as well as many other specifications, such as the Smooth Ambiguity Preferences model. The general *Uncertainty-Averse Preferences* representation is

$$W(f(\cdot)) \equiv \min_{\mu \in \Delta(\mathcal{S})} G\left(\int U(f(\cdot))d\mu, \mu\right) \tag{13.14}$$

where the function $G : \mathbb{R} \times \Delta(\mathcal{S}) \rightarrow (-\infty, +\infty]$ is quasiconvex, increasing in the first variable, and satisfies $\inf_{\mu \in \Delta(\mathcal{S})} G(t, \mu) = t$ for all t .

Due to their generality, *Uncertainty-Averse Preferences* representations are only unique in a weak sense. For a fixed utility function, there is a minimal function $G(\cdot, \cdot)$. If a positive affine transformation of U is taken, G must be suitably rescaled.

MEU preferences are, of course, also *Uncertainty-Averse Preferences*, with

$$G(t, \mu) = \begin{cases} t & \mu \in C \\ \infty & \mu \notin C \end{cases} .$$

Variational Preferences correspond to the special case

$$G(t, \mu) = t + c(\mu)$$

whereas Confidence-Function Preferences obtain if

$$G(t, \mu) = \begin{cases} \frac{t}{\varphi(\mu)} & \varphi(\mu) \geq \alpha \\ \infty & \varphi(\mu) < \alpha \end{cases} .$$

Finally, recall that Smooth Ambiguity Preferences also satisfy *Uncertainty Aversion*, provided the second-order utility function $\varphi(\cdot)$ is concave. Under additional technical assumptions, Cerreia-Vioglio et al. (2011a) show that the Smooth Ambiguity Preference representation can be written as in equation (13.14), with

$$G(t, \mu) = t + \min_{M' \in \Gamma(\mu)} I_t(M'|M)$$

where M' is the second-order probability measure in the Smooth Ambiguity Preferences representation, $\Gamma(\mu)$ is the set of second-order measures $N' \in \Delta(\Delta(\mathcal{S}))$ that reduce to μ , in the sense that

$$\int_{\Delta(\mathcal{S})} \pi(E)dM'(\pi) = \mu(E)$$

for every event $E \subseteq \mathcal{S}$, and $I_t(M'|M)$ is a so-called *statistical distance function*, i.e., a nonnegative, lower semicontinuous, quasiconvex function that satisfies $I_t(M|M) = 0$ (a specific formula is provided by Cerreia-Vioglio et al., 2011a). In other words, Nature attempts to minimize the expected utility of the act under consideration by choosing a probability distribution μ , at a cost given by the result of the minimization problem above. The latter has an interesting “prior uncertainty” interpretation: out of all second-order probability measures M' whose reduction equals μ , Nature chooses the one closest to M . As for Multiplier Preferences, the intuition is that, while M is the individual’s “preferred” second-order belief, measures M' close to it are also plausible.

13.5.10 Gul and Pesendorfer’s Expected Uncertain Utility

Gul and Pesendorfer (2013) introduce a new model that can accommodate Ellsberg preferences, as well as source preference. Their model applies to a setting *a la* Savage, with an infinite state space \mathcal{S} and outcomes in some real interval $[l, m]$. The key intuition is that the individual is able to assign probabilities to a (suitably rich) collection of events, but displays complete ignorance with respect to all other events.

Formally, the proposed representation, *Expected Uncertain Utility (EUU)* starts with a σ -algebra Σ , which comprises the events whose likelihood the individual can confidently assess. A countably additive and nonatomic measure μ represents beliefs on Σ . The *EUU* model ranks Σ -measurable acts according to expected utility; these acts are “unambiguous.” The question is how to rank an arbitrary Savage act $f: \mathcal{S} \rightarrow [l, m]$ that is not Σ -measurable. To do so, according to *EUU*, the individual *brackets* any act f between a lower Σ -measurable approximation f_1 and an upper Σ -measurable approximation f_2 , in the sense that

$$\mu(\{s: f_1(s) \leq f(s) \leq f_2(s)\}) = 1$$

and furthermore such pair (f_1, f_2) is the “tightest” approximation of f : if there is another pair (g_1, g_2) with the same property, then $\mu(\{s: g_1(s) \leq f_1(s) \leq f_2(s) \leq g_2(s)\}) = 1$. The last ingredient of the *EUU* representation is a continuous, increasing *interval utility function*, that is, a map $U(\cdot, \cdot)$ from the set $I = \{(x, y): x \leq y\}$ to \mathbb{R} . The interpretation is that $U(x, y)$ is the utility of receiving a prize somewhere in the interval $[x, y]$, when the individual is completely ignorant as to which prize in this interval he or she will actually receive. The representing function is then

$$W(f(\cdot)) \equiv \int U(f_1(\cdot), f_2(\cdot)) d\mu. \quad (13.15)$$

To see how this representation works, consider a bet on an event A , i.e., an act $f = (x, A; y, \sim A)$ with $x > y$ (which implies that x is strictly preferred to y). If the event

A is not in Σ , the measure μ does not assign a probability to it; however, we can consider the quantity $\sup \{\mu(E) : E \in \Sigma, E \subseteq A\}$. There will be an event $E_1 \in \Sigma$ that attains this supremum. Similarly, there is an event $E_2 \in \Sigma$ that attains $\sup \{\mu(E) : E \in \Sigma, E \subseteq \sim A\}$. A bracketing of the binary act f is then given by $f_1 = (x, E_1; y, \sim E_1)$ and $f_2(x, \sim E_2; y, E_2)$. Then

$$W(f(\cdot)) = U(x, x) \cdot \mu(E_1) + U(y, y) \cdot \mu(E_2) + U(y, x) \cdot \mu(\sim (E_1 \cup E_2)).$$

This has an intuitive interpretation. On the unambiguous event $E_1 \subseteq A$, the individual receives x for sure. Similarly, he or she receives y on $E_2 \subseteq \sim A$. How about states that belong to neither E_1 nor E_2 ? Some of these states are in A , and some are in $\sim A$, but the individual is unable to rank the relative likelihood of $A \setminus E_1$ and $(\sim A) \setminus E_2$. Thus, according to *EUU*, he or she assigns to all such states the interval utility value $U(y, x)$ —by definition, the utility assigned to receiving some prize between y and x .

The authors provide an elegant, Savage-style axiomatization of *EUU*. The events in Σ are identified as those events for which Savage’s Sure-Thing Principle P2 holds, provided it holds for their complement as well. An axiom in the spirit of Savage’s P6 ensures the existence of a rich collection of such events. Furthermore, a rich collection of events similar to $\sim(E_1 \cup E_2)$ in the above example is also assumed. The reader is referred to the original paper for further details.

13.5.11 Bewley’s Incomplete-Preference Model

A common feature of all the preceding models is that, despite the presence of ambiguity, the individual was assumed to be able to rank all acts, or all suitably measurable acts. Bewley (2002, 2011) proposes a different approach: he assumes that, due to ambiguity, some acts may simply not be comparable. For example, Bewley’s take on the Three-Color Paradox is that, quite simply, the individual is unable to compare a bet on red to a bet on black, or a bet on red or yellow to a bet on black or yellow.

He proposes a representation of incomplete preferences that features a set \mathcal{C} of probability distributions, as in *MEU*. According to his model, for any two acts f, g , f is preferred to g if and only if

$$\int U(f(\cdot)) d\mu \geq \int U(g(\cdot)) d\mu \quad \forall \mu \in \mathcal{C}.$$

This is sometimes called the *unanimity rule*: all priors in \mathcal{C} must “agree” that f is better than g . This representation plays a role in recent developments on the definition of ambiguous beliefs, as in Section 13.6.

To complement this incomplete decision criterion, Bewley suggests an *inertia assumption*, which is essentially a form of status-quo bias: if an act f describes the individual’s

initial endowment, or default option, and there is no act g that is strictly better than f according to the above criterion, the individual retains f .

The connection between inertia and the status-quo bias is formalized in Ortoleva (2010). This paper takes as primitive a collection of preferences over acts, of which one is an “unconditional” preference, and the others are preferences “conditional” upon an exogenously given status-quo act. Under suitable axioms, the joint representation of these preferences identifies an underlying incomplete, Bewley-style preference, and involves inertia.

13.5.12 Models with “Objective Ambiguity”

Ambiguity and ambiguity attitudes have also been modeled by extending von Neumann and Morgenstern’s analysis of expected utility under objective uncertainty rather than Savage’s axiomatization of expected utility under subjective uncertainty.

One approach in this strand of the literature identifies ambiguous bets and acts with *sets of lotteries*. For instance, a bet on the ambiguous color black in the Three-Color Paradox can be identified with the set $\{(\$100, \alpha; \$0, 1-\alpha) : \alpha \in [0, 2/3]\}$. An unambiguous act such as a bet on red is instead identified with the singleton set $\{(\$100, 1/3; \$0, 2/3)\}$. These sets of lotteries are *not* part of a parametric representation: in this model, they are the objects of choice. A state space is *not* specified. The individual is characterized by a preference ordering \succsim over sets of objective lotteries. Preferences over singleton sets obey the von Neumann–Morgenstern axioms; the interest is chiefly in the way non-singleton sets are evaluated.

Olszewski (2007) characterizes the α -MEU criterion, with α strictly between 0 and 1. Ahn (2008) proposes the criterion

$$V(A) = \frac{\int \phi(U(\mathbf{P})) d\mu}{\mu(A)}$$

where μ is a measure over all lotteries that is “conditioned” on each menu A , U is a von Neumann–Morgenstern utility, and ϕ is an increasing function. Dumav and Stinchcombe (2013) provide a general analysis of linear representations of preferences over sets of lotteries, and investigate the connection with models that do employ a state space, such as the ones described in Sections 13.5.1–13.5.11.

A related literature considers preferences over tuples consisting of a set of probability distributions on a state space, an act, and possibly a reference probability. Sets of probability distributions can be interpreted as representing *imprecise information*, perhaps obtained by econometric techniques. The issue of interest is how to model the individual’s attitudes toward such imprecise information.

Gajdos et al. (2008) consider preferences over pairs (P, f) of probability sets and acts (note that it is assumed that comparisons with different sets of probability distributions

are also observed). A first, general representation evaluates such a pair by a maxmin criterion applied to a suitably selected subset of P . Under additional assumption, an elegant representation is obtained for finite state spaces, in which (P, f) is evaluated according to a weighted average of the minimum expected utility over P and the expected utility with respect to the *Steiner point* of P (a construct which generalizes the notion of center of gravity). Wang (2003) instead considers preferences over tuples (P, p, f) , where P is a set of probability distributions, p is a reference prior (a distinguished element of P) and f is an act. He axiomatizes a *MEU* criterion in this set, as well as an analog of the Multiplier Preferences representation of Hansen and Sargent (2001).

Viscusi (1989) proposes a different approach within the von Neumann–Morgenstern framework, termed *Prospective Reference Theory*. He suggests that “stated” probabilities in a lottery are not necessarily taken at face value by the individual. Rather, they are interpreted as signals about the “true” probabilities. Individuals then behave according to a weighted average of the stated probability and a prior, or baseline, probability; importantly, they may attribute different weights to the stated probabilities of different events. In particular, they may assign greater weight to the stated probability of a red draw from the Three-Color Ellsberg urn than to the stated probability of a black draw; correspondingly, they may assign a greater weight to the stated probability of a yellow or black draw than to the stated probability of a red or black draw. This allows the Prospective Reference Theory model to match the modal preferences in the Three-Color Paradox.

13.6 RECENT DEFINITIONS AND EXAMPLES

13.6.1 Recent Definitions of Ambiguity and Ambiguous Events

The parametric representations of preferences discussed in Section 13.5 suggest conditions under which an event E may be deemed “ambiguous.” For instance:

- The intuitive interpretation of the Maxmin Expected Utility model is that the individual may be unable to assign a unique probability to certain events; thus, an event E may be called “ambiguous” if $\min_{\mu \in \mathcal{C}} \mu(E) < \max_{\mu \in \mathcal{C}} \mu(E)$.
- A similar interpretation applies to the Smooth Ambiguity Preferences model: E is ambiguous if $\min_{\mu \in \text{supp} M} \mu(E) < \max_{\mu \in \text{supp} M} \mu(E)$, where M is the second-order prior and “supp” denotes the support of a measure.
- The suggested interpretation of the Choquet Expected Utility model is that the individual expresses his or her lack of confidence in likelihood assessments by violating additivity. Thus, E may be called “ambiguous” if the capacity ν satisfies $\nu(E) + \nu(\sim E) < 1$.
- In the Vector Expected Utility model, the adjustment function captures departures from expected utility, so E may be called “ambiguous” if $A(\text{cov}(1_E, \zeta_1), \dots, \text{cov}(1_E, \zeta_n)) \neq 0$.

However, these definitions are all closely tied to specific functional forms. Several authors have proposed definitions of ambiguous and unambiguous events that are stated solely in terms of a decision-maker's preferences. We now review two that are, in a sense, polar opposites.

13.6.1.1 The Epstein-Zhang (2001) Definition

Epstein and Zhang (2001) propose a definition of ambiguous and unambiguous events that generalizes the preference pattern in the Three-Color Paradox. Recall that, loosely speaking, the modal preferences in that paradox exhibit a kind of *preference reversal*: the individual prefers to bet on red rather than on black, but once yellow is added to the event on which the favorable prize \$100 is delivered, the decision maker prefers to bet on black or yellow rather than on red or yellow. As noted in Section 13.5.2, this indicates that there is a “complementarity” between yellow and black, but not between red and black. The basic idea behind the approach is to take the fact that yellow is “complementary” to another color (in this case, black) as the *defining feature* of an ambiguous event. More generally, an event T is an *unambiguous* event if it is not “complementary” to some other, disjoint event, in the sense that a preference reversal similar to that in the Three-Color Paradox never obtains.

The Epstein-Zhang definition adds two “robustness checks” to the above intuition. To motivate the first, consider a five-color version of the Three-Color Paradox, and allow for a set of prizes that contains \$0 and \$100, but other elements as well. The acts under consideration are

	30 balls red	60 balls black yellow		60 balls white green	
a_1	\$100	\$0	\$0	w'	w''
a_2	\$0	\$100	\$0	w'	w''
a_3	\$100	\$0	\$100	w'	w''
a_4	\$0	\$100	\$100	w'	w''

We have added two colors, white and green, and assigned arbitrary prizes to all four acts at the corresponding states; importantly though, the prize in case a white ball is drawn is the same (namely w' for all four acts, and similarly (w'') if a green ball is drawn. Regardless of information available regarding the number of white or green balls, the intuition in the Three-Color Paradox should still apply: yellow should still be “complementary” to black, regardless of the number of white or green balls. Conversely, if we observed the preference pattern $a_1 \succ a_2$ and $a_3 \prec a_4$ for *some* choice of prizes w' , w'' , we should still conclude that yellow is an ambiguous color.

For the second robustness check, consider the following variation on the previous one:

	30 balls	60 balls		60 balls	
	red	black	yellow	white	green
a_1	\$100	\$0	z'	w'	w''
a_2	\$0	\$100	z'	w'	w''
a_3	\$100	\$0	z'	w'	w''
a_4	\$0	\$100	z'	w'	w''

The prizes \$0 and \$100 have been replaced with arbitrary prizes z and z' respectively. One may argue that if z is close to \$0 and z' is close to \$100, the rankings $a_1 \succ a_2$ and $a_3 \prec a_4$ should still be plausible. Conversely, if we observed such rankings, we might say that the individual likes acts that deliver *similarly good* prizes in case of a black or yellow draw: this, too is a form of complementarity between yellow and black, so we can take it as indication that yellow is an ambiguous color. Indeed, Epstein and Zhang suggest that this conclusion would be warranted so long as a preference reversal occurs for *some* choice of z and z' , even if it does not obtain with $z = \$0$ and $z' = \$100$. Accordingly, an event T is defined to be *unambiguous* if (a) for all disjoint subevents A, B of $\sim T$, all prizes x, y, z, z' with $x \succ y$ and all acts h ,

$$(x, A; y, B; z, T; h, (\sim T) \setminus (A \cup B)) \succcurlyeq (y, A; x, B; z, T; h, (\sim T) \setminus (A \cup B))$$

implies

$$(x, A; y, B; z', T; h, (\sim T) \setminus (A \cup B)) \succcurlyeq (y, A; x, B; z', T; h, (\sim T) \setminus (A \cup B))$$

and furthermore (b) the same is true if T is replaced everywhere with $\sim T$. An event is *ambiguous* if it is not unambiguous.

Epstein and Zhang observe that, in general, the collection of unambiguous events is *not* an algebra. In particular, the intersection of two unambiguous events may be ambiguous. Zhang (2002) proposes the following examples: an urn contains white, red, green and blue balls. It is known that there are 50 red and blue balls, and 50 green and blue balls; the total number of balls is 100. Intuitively, the events “red or blue” and “green or blue” are unambiguous, but their intersection, “blue,” is ambiguous.⁶⁵ However, under suitable assumptions (see below), the collection of unambiguous events contains the state space \mathcal{S} and is closed under complements and countable *disjoint* unions; in other words, it is a λ -system.

Epstein and Zhang consider preferences over Savage acts, and assume that preferences satisfy Savage’s postulate P3 (Monotonicity). They then consider the family of *unambiguous acts*, that is, acts measurable with respect to unambiguous events, and show that, under suitable axioms, there exists a unique, convex-ranged and countably additive

⁶⁵ To clarify, this is *not* a consequence of the above definition, because we have not specified preferences. However, if we did so (for instance, by considering *MEU* preferences with $\mathcal{C} = \{p \in \Delta(\mathcal{S}) : p(s_r, s_b) = p(s_g, s_b)\}$ in the obvious notation), the definition would indeed lead to the intuitive conclusions in the text.

probability measure μ such that preferences over unambiguous acts are *probabilistically sophisticated* with respect to μ .

The interpretation is as follows: Recall that the Ellsberg Paradoxes contradict not just expected utility theory, but the more general assumption of probabilistic sophistication. Conversely, a probabilistically sophisticated individual does not exhibit Ellsberg preferences. Since the definition of unambiguous events was *motivated* by analogy with the Three-Color Paradox, it is almost a tautology that preferences over acts measurable with respect to unambiguous events cannot exhibit Ellsberg-type behavior. But the Epstein-Zhang result says much more: preferences over such acts uniquely pin down a probability measure over unambiguous events, and indeed they are probabilistically sophisticated.

If the decision-maker is probabilistically sophisticated in the sense of Machina and Schmeidler (1992), then the above definition deems *all* events to be unambiguous, and the Epstein-Zhang axioms hold; thus, their result is consistent with Machina and Schmeidler's. The Epstein-Zhang result is mainly of interest when the overall preferences are *not* probabilistically sophisticated; in such cases, it identifies a subset of events and preferences for which probabilistic sophistication holds.

Underlying the terminology employed by Epstein and Zhang is the implicit assumption that *ambiguity means a violation of probabilistic sophistication*. As we shall see, the approach we discuss next makes a different implicit assumption.

13.6.1.2 The Ghirardato-Maccheroni-Marinacci/Nehring Definition

An alternative approach has been suggested by Ghirardato et al. (2004) and Nehring (2001). We follow the exposition in the former paper.

Ghirardato et al. consider preferences over horse-roulette acts. Their starting point is the observation that Ellsberg preferences violate the Independence axiom, as noted in Section 13.5. They then propose to examine *unambiguous* preference rankings—that is, rankings that are preserved by *all* mixtures, or equivalently, for which there is no preference reversal. Formally, they define the unambiguous preference relation \succ^* by

$$f \succ^* g \quad \text{iff} \quad \forall \lambda \in (0, 1], \quad \forall h : \lambda \cdot f + (1 - \lambda) \cdot h \succ \lambda \cdot g + (1 - \lambda) \cdot h. \quad (13.16)$$

An equivalent definition is that \succ^* is the maximal subrelation of \succ that satisfies Independence.

If the primitive preference relation \succ satisfies a minimal set of axioms (Monotonicity, Continuity and *Risk Independence*—that is, the Independence axiom restricted to constant acts only), it will follow that \succ^* is a *Bewley preference* (Section 13.5.11).⁶⁶ Therefore, there is a set \mathcal{C} of probability measures, or *relevant priors*, such that, for all acts $f, g, f \succ^* g$

⁶⁶ Ghirardato et al. show that this is the case under the stronger assumption of Certainty Independence. Siniscalchi (2009a, Lemma 2) proves it under the weaker assumption of Weak Certainty Independence. Cerreia-Vioglio et al. (2011) establish the general result cited in the text.

if and only if $\int U(f(\cdot))d\mu \geq \int U(g(\cdot))d\mu$ for all $\mu \in \mathcal{C}$, where U is the utility function that represents the primitive preference relation when restricted to constants.

Ghirardato et al. call an act *crisp* if it is unambiguously indifferent to a constant roulette lottery: that is, there is a prize (roulette lottery) x such that $f \sim^* x$.⁶⁷ The intuition is straightforward: preferences over roulette lotteries are consistent with expected utility, and hence not subject to ambiguity; thus, an act that is treated just like a lottery, even when mixing with other acts, is also not affected by ambiguity. It can be shown that an act is crisp if and only if its expected utility $\int U(f(\cdot))d\mu$ is the same for all relevant priors $\mu \in \mathcal{C}$. Thus, effectively, preferences over crisp acts are consistent with expected utility.

Finally, an event E is deemed *unambiguous* if every bet on E , that is, any binary act of the form $(x, E; y, \sim E)$ with $x \succ y$, is crisp. From the characterization of crisp acts, it is immediate that this is equivalent to the property that $\mu(E)$ is constant for all relevant priors $\mu \in \mathcal{C}$: that is, the probability measures in the set \mathcal{C} “collapse” at E . Conversely, E is ambiguous if and only if $\min_{\mu \in \mathcal{C}} \mu(E) < \max_{\mu \in \mathcal{C}} \mu(E)$. Recall that this is the “intuitive” notion of ambiguity for events in the *MEU* representation. However, in the approach advocated by Ghirardato et al./Nehring, this notion applies to a broad class of preferences that includes virtually all the models of Section 13.5, and it is fully characterized by a behavioral condition (the assertion that bets on E are crisp).

There are several formal differences with the Epstein-Zhang approach. Epstein and Zhang adopt the Savage framework, start from a definition of unambiguous events, and derive a definition and characterization of unambiguous acts. Ghirardato et al. adopt the horse-roulette act framework, start from a definition of crisp (i.e., unambiguous) acts, and derive a characterization of unambiguous events. Conceptually, however, the key difference is that Epstein and Zhang regard any probabilistically sophisticated preference relation (or restriction thereof) as unambiguous; Ghirardato et al./Nehring, on the other hand, only deem expected utility preferences to be unambiguous. Both approaches classify Ellsberg preferences as ambiguity-sensitive, and expected utility preferences as unambiguous, but the Epstein-Zhang notion of ambiguity is seen to be more demanding.

For instance, consider a *MEU* preference on acts defined over a three-point state space, with $\mathcal{C} = \{p : p(s) \geq 1/9 \text{ for all } s\}$. These preferences can be equivalently described as rank-dependent, with distortion function $\delta(x) = x^2$ and a uniform prior. According to Ghirardato et al./Nehring, the only unambiguous events are the empty set and the entire space; Epstein and Zhang consider all events unambiguous. Ultimately, it may be difficult to decide which interpretation is correct.

One advantage of the Ghirardato et al./Nehring approach is that computing the set \mathcal{C} of priors in the Bewley representation of unambiguous preferences, and hence

⁶⁷ Ghirardato et al. actually use a different definition that relies on Certainty Independence. We use here the more general definition in Cerreia-Vioglio et al. (2011).

determining which events are unambiguous, is relatively straightforward for many parametric models of preferences. For preferences that satisfy Certainty Independence (e.g., *CEU* and *MEU*), Ghirardato et al. show that \mathcal{C} is the Clarke (1983) differential of the function that represents preferences, evaluated at 0. For instance, an expected utility preference with prior μ has a single relevant prior, namely μ —a basic check that indicates the soundness of the definition; a *MEU* preference with priors \mathcal{C} has precisely \mathcal{C} as the set of relevant priors. Ghirardato and Siniscalchi (2012) provide a general characterization that does not require Certainty Independence, and thus covers virtually all the models in Section 13.5 (see Ghirardato and Siniscalchi (2010) for explicit calculations of the relevant priors for Variational, Multiplier, Uncertainty-Averse, Smooth Ambiguity and *VEU* Preferences).

Ghirardato and Siniscalchi (2012) also propose a local notion of unambiguous preference, provide a differential characterization, and relate it to the global definition of unambiguous preference described above. They also relate local unambiguous preferences to optimizing behavior.

13.6.2 Recent Definitions of Ambiguity Aversion

The Uncertainty Aversion Axiom of Schmeidler (1989) represents the first attempt to formalize the notion that individuals dislike ambiguity. As discussed in Section 13.5, the intuition is that, by mixing two acts, the individual may be able to hedge against variation in utilities, much like, by forming a portfolio consisting of two or more assets, one can hedge against variation in monetary payoffs.

However, other attempts to characterize a dislike for ambiguity have been proposed in the literature.

13.6.2.1 Epstein's Comparative Definition

Epstein (1999) critiques the Uncertainty Aversion Axiom, focusing on the case of *CEU* preferences. Recall from Section 13.5.2 that a *CEU* preference satisfies Uncertainty Aversion if and only if the representing capacity ν is convex: that is, if and only if, for every pair of events E, F , $\nu(E \cup F) + \nu(E \cap F) \geq \nu(E) + \nu(F)$. Epstein (p. 582) observes that this condition is “neither necessary nor sufficient” for *CEU* preferences to reproduce the modal choices in the Three-Color Paradox. For example, if $\nu(\{s_r\}) = 8/24$, $\nu(\{s_b\}) = \nu(\{s_y\}) = 7/24$, $\nu(\{s_b, s_y\}) = 13/24$ and $\nu(\{s_r, s_y\}) = \nu(\{s_r, s_b\}) = 1/2$, then the modal rankings $a_1 > a_2$ and $a_4 > a_3$ are obtained, but the capacity thus defined is not convex: for instance, $13/24 = \nu(\{s_b, s_y\}) = \nu(\{s_b\} \cup \{s_y\}) + \nu(\{s_b\} \cap \{s_y\}) < \nu(\{s_b\}) + \nu(\{s_y\}) = 14/24$. Accordingly, convexity of the capacity is not necessary. Suppose instead that $\nu(\{s_r\}) = 1/12$, $\nu(\{s_b\}) = \nu(\{s_y\}) = 1/6$, $\nu(\{s_b, s_y\}) = 1/3$ and $\nu(\{s_r, s_y\}) = \nu(\{s_r, s_b\}) = 1/2$; this capacity is convex, but the corresponding *CEU* preference generates the uncertainty-loving rankings $a_2 > a_1$ and $a_3 > a_4$! This shows that convexity of the capacity is not sufficient—it does not even prevent uncertainty-loving patterns of behavior.

The fact that the Uncertainty Aversion Axiom is neither necessary nor sufficient for the *CEU* model to exhibit Ellsberg Preferences naturally casts doubts on the implications for this axiom in other decision models. With this motivation, Epstein sets out to provide an alternative definition of ambiguity aversion. The approach introduced in his 1999 paper is refined and, in a sense, brought to completion in Epstein and Zhang (2001) as discussed in Section 13.6.1. We thus follow the exposition in the latter paper.

The decision setting is that of Savage, so in particular, acts map states to prizes rather than lotteries over prizes. The starting point of Epstein’s analysis is Yaari’s (1969) notion of *comparative risk aversion*, namely that a preference relation \succsim^2 is *more risk-averse* than another preference relation \succsim^1 if, for all acts h and prizes x

$$h \succsim^2 x \Rightarrow h \succsim^1 x \quad \text{and} \quad h \succ^2 x \Rightarrow h \succ^1 x. \tag{13.17}$$

That is, whenever \succsim^2 prefers an (intuitively, risky) act h to a (riskless) prize x , so does \succsim^1 ; this leaves open the possibility that some act h' , which \succsim^2 may rank inferior to a prize x' , may instead be ranked above x' by \succsim^1 . With monetary prizes and expected utility preferences, this definition is equivalent to the usual characterizations in terms of the relative concavity of the von Neumann–Morgenstern utility functions. (The advantage of this definition is that it applies to non-*EU* models just as well, and hence provides a sort of “neutral testing ground” for comparing risk attitudes.)

To define a notion of risk aversion *per se*, as opposed to comparative risk aversion, one needs to fix a *benchmark*—that is, to decide which preferences to declare “risk-neutral.” With monetary prizes, we can take this benchmark or reference to be *expected-value preferences*. In the general case, we say that a preference \succsim is risk-averse if it is more risk-averse than *some* expected-value preference \succsim^m . Again, with expected utility preferences, this coincides with the usual characterization in terms of concavity of the utility function; however, the same definition can be usefully applied to non-*EU* preferences as well.

Epstein’s suggestion is to define ambiguity aversion in an analogous fashion. Recall that Epstein and Zhang (2001) identify a class (a λ -system) of unambiguous events: acts measurable with respect to this class are thus unambiguous acts. First, say that \succsim^2 is *more ambiguity-averse* than \succsim^1 if, for every arbitrary act f and unambiguous act h ,

$$f \succsim^2 h \Rightarrow f \succsim^1 h \quad \text{and} \quad f \succ^2 h \Rightarrow f \succ^1 h. \tag{13.18}$$

As a consequence of equation (13.18), \succsim^2 and \succsim^1 rank unambiguous acts the same way. However, \succsim^1 may prefer some arbitrary act f to an unambiguous act h , when \succsim^2 instead prefers the unambiguous act; on the other hand, the opposite never happens. In this sense, \succsim^2 is more ambiguity-averse than \succsim^1 .

To define ambiguity aversion *per se*, a benchmark must be selected. Consistently with the approach discussed in Section 13.6.1, Epstein and Zhang adopt probabilistically sophisticated preferences as the benchmark. Thus, they propose that *a preference relation is ambiguity-averse if it is more ambiguity-averse than some probabilistically sophisticated preference relation.*

Epstein and Zhang (2001, Corollary 7.4) characterize *CEU* preferences that are ambiguity-averse in this sense. For *MEU* preferences, Epstein (1999, Thm. 3.3) shows that, if all priors in the set \mathcal{C} coincide on unambiguous events, then ambiguity aversion holds (the next subsection indicates why a qualification is required). Epstein (1999, Sect. 3.4) shows that the Ellsberg choices in the Three-Color Paradox imply ambiguity aversion, and that, conversely, ambiguity aversion rules out the behavioral pattern that is intuitively associated with an attraction to ambiguity.

13.6.2.2 Ghirardato and Marinacci's Comparative Definition

An alternative approach is proposed by Ghirardato and Marinacci (2002), and then further developed by Ghirardato et al. (2004) and Cerreia-Vioglio et al. (2011b). Ghirardato and Marinacci consider both the Savage setting, with acts mapping to prizes, and the horse-roulette act framework, with acts mapping to objective probability distributions over prizes. They restrict attention to preferences that admit a *CEU* representation on *binary* acts, but are otherwise arbitrary (for instance, all *MEU* preferences satisfy this restriction). In the Savage environment, they further restrict comparisons to preferences that satisfy a *cardinal symmetry* requirement, which implies that a common von Neumann-Morgenstern utility can be used in the *CEU* representation on binary acts.

The starting point is again Yaari's (1969) approach, but these authors focus on comparisons between acts and constant prizes. Formally, Ghirardato and Marinacci deem a preference \succsim^2 *more ambiguity-averse* than \succsim^1 if, for all acts f and prizes x ,

$$f \succsim^2 x \Rightarrow f \succsim^1 x \quad \text{and} \quad f \succ^2 x \Rightarrow f \succ^1 x. \quad (13.19)$$

The interpretation is analogous to that provided in the previous subsection for Epstein's definition: whenever \succsim^2 prefers an (intuitively, potentially ambiguous) act f to a (clearly unambiguous) constant x , so does \succsim^1 ; the opposite, however, need not hold.

To define ambiguity aversion per se, Ghirardato and Marinacci take as benchmark *expected utility* preferences. Thus, according to these authors, *a preference is ambiguity-averse if it is more ambiguity averse than some EU preference.*

Cerreia-Vioglio et al. (2011b) provide a characterization for a general class of preferences that satisfy Monotonicity, Continuity and Risk independence in a horse-roulette act framework (or in environments with subjective mixtures, as discussed at the beginning of Section 13.5). This is also the class of preferences for which a Bewley characterization of unambiguous preferences can be provided (Section 13.5.11). Any such preference relation is ambiguity-averse in the sense of Ghirardato and Marinacci if and only if there exists a probability measure μ such that

$$\int U(f(\cdot)) d\mu \geq V(f(\cdot))$$

where $V(\cdot)$ as usual denotes the utility index assigned to acts. In other words, there must exist an expected utility representation, with utility function U , which dominates $V(\cdot)$ pointwise. Some immediate consequences of this characterization (and of earlier ones provided by Ghirardato and Marinacci (2002) and Ghirardato, et al.) are as follows:

- All Maxmin Expected Utility preferences are ambiguity-averse.
- Choquet Expected Utility preferences are ambiguity-averse if and only if the capacity ν has a *non-empty core*: that is, if there exists a probability measure p such that $p(E) \geq \nu(E)$ for every event E .
- Vector Expected Utility preferences are ambiguity-averse if and only if the adjustment function A is nonpositive.
- Smooth Ambiguity Preferences with a concave second-order utility function φ are ambiguity averse.

As will be seen in Section 13.6.3 (in particular footnote 69), there are interesting and experimentally verified preference patterns that are inconsistent with Schmeidler's Uncertainty Aversion Axiom, but are consistent with Ghirardato and Marinacci's notion of ambiguity aversion. On the other hand, Cerreia-Vioglio et al. (2011b, Example 2) show that Schmeidler's Uncertainty Aversion Axiom is in general not stronger than Ghirardato and Marinacci's notion of ambiguity aversion.

A comparison with the Epstein/Epstein-Zhang approach is in order. The main conceptual difference is that—as was the case for the Epstein vs. Ghirardato et al./Nehring notions of ambiguity—the choice of the ambiguity-neutral benchmark is different. Epstein adopts probabilistic sophistication; Ghirardato and Marinacci choose expected utility. Epstein (1999) observes that using *EU* as benchmark implies that, for instance, the modal preferences in Savage's restatement of the Allais paradox (Savage (1954, p.103)) are classified as ambiguity averse. There is, of course, nothing that evokes ambiguity or ambiguity attitudes in the Allais bets, even when translated in a Savage framework: the composition of the urn is known to the individual, who presumably trusts the experimenter, etc. But the force of this argument derives in part from the particular example it considers. One can easily imagine choice situations in which ruling out ambiguity aversion on the basis of available information is not straightforward. Suppose for instance that an individual is asked to bet on which of three cities in an unfamiliar foreign country has the highest average temperature in the month of May. Suppose that this individual's preferences admit the *MEU* representation of Section 13.5.1, with $\mathcal{C} = \{p : p(s) \geq 1/9 \text{ for all } s\}$. We noted above that these preferences can equivalently be described as rank-dependent (*CEU*), with distortion function $\delta(x) = x^2$ and a uniform prior. Since these preferences are probabilistically sophisticated, Epstein classifies them as ambiguity-neutral; Ghirardato and Marinacci instead classify them as ambiguity averse.

This alternative example does not provide clear indications as to whether or not ambiguity plays a role.⁶⁸

The Epstein definition is accordingly more “conservative”: if a pattern of behavior is deemed “ambiguity-averse,” it *surely* has features that are inconsistent with probabilistic beliefs. Thus, for instance, if one was interested in demonstrating the presence and relevance of ambiguity in a given economic environment, applying the Epstein definition would provide more convincing evidence. On the other hand, the Ghirardato–Marinacci definition captures a broader range of behavioral patterns that *may* be influenced by ambiguity, and is easily characterized in a variety of parametric preference models.

13.6.2.3 Smooth Ambiguity Aversion

As anticipated in Section 13.5.4, Klibanoff et al. (2005) propose a notion of ambiguity aversion that is specific to the Smooth Ambiguity Preferences model.

These authors identify an individual with a collection $\{\succsim_K, \succsim_K^{(2)}\}$ of first- and second-order preferences, indexed by the *support* K of the second-order probability measure M representing each pair. They assume that all first-order preferences in any such collection exhibit the same risk attitudes, and hence pin down the same utility $U(\cdot)$; similarly, they assume that risk preferences are the same for all second-order preferences, which pins down the utility function $\varphi(\cdot)$ in the *EU* representation of second-order acts, and consequently also the second-order utility $\varphi(\cdot)$. The objective is to provide a condition on all elements of the collection $\{\succsim_K, \succsim_K^{(2)}\}$ that characterizes second-order risk aversion, i.e., concavity of the function $\varphi(\cdot)$.

Their construction is as follows: Recall that any act f and any first-order probability measure μ induce a lottery over prizes, $(\dots; x_i, \mu(f^{-1}(x_i)); \dots)$. This lottery has expected utility $\int U(f(\cdot))d\mu$. Since the second-order measure M is defined over first-order probabilities μ , it induces a *lottery over expected utilities*, which in the case of a finite support $K = \{\dots, \mu_j, \dots\}$ can be expressed as

$$\mathbf{P}_f^* \equiv \left(\dots; \int U(f(\cdot))d\mu_j, M(\mu_j); \dots \right).$$

The smooth ambiguity evaluation of the act f corresponds to computing the “expected utility” of the lottery \mathbf{P}_f^* with respect to the second-order utility $\varphi(\cdot)$. Hence, to characterize concavity of the latter, we need to compare f with a prize $x_f(M)$ whose utility is precisely the expected value of \mathbf{P}_f^* —equivalently, the *average* expected utility of f . Such a prize exists because U has a convex range. Klibanoff et al. then show that the

⁶⁸ Epstein and Zhang’s characterization of ambiguous and unambiguous events, and hence their definition of ambiguity aversion, requires a rich state space. However, Epstein’s original definition does not restrict the state space, and classifies any probabilistically sophisticated preference as ambiguity neutral. Thus, our discussion of the temperature example is indeed consistent with Epstein (1999).

second-order utility $\varphi(\cdot)$ is concave if and only if the individual is *smooth ambiguity averse*, in the sense that for every support K with associated second-order measure M , and for every act f , the individual prefers the prize $x_f(M)$ to f — formally, $x_f(M) \succ_K f$. A similar construction can be used to characterize comparative smooth ambiguity aversion.

It should be noted that, due to the way $x_f(M)$ is defined, smooth ambiguity aversion depends crucially on the representation of second-order preferences. Thus, the qualifications on second-order preferences, and on the interpretation of the smooth representation noted in Section 5.4, similarly apply here. Furthermore, again because its definition involves elements of the functional representation, smooth ambiguity aversion cannot be directly compared with the Epstein and Ghirardato–Marinacci definitions (though of course one can ask whether specific parameterizations satisfy one of these definitions, as we did in the previous section).

13.6.3 Recent Examples of Ambiguity Aversion

The models of ambiguity aversion reviewed in Section 13.5 were designed to capture aspects of ambiguity aversion revealed in Ellsberg’s original examples, each of which only involved two possible outcome values (Ellsberg used \$0 and \$100). However, additional aspects of ambiguity aversion, which can only reveal themselves in choices involving three or more outcome values, have since come to light, and pose challenges to some of these models.

A simple example of this is illustrated in the following table, based on an urn containing 50 red or black balls, and 51 green or yellow balls.⁶⁹ In the first pair of bets, it is likely that ambiguity aversion would outweigh bet a_1 ’s slight 50:51 disadvantage, causing it to be preferred over a_2 . In the second pair, the difference in the bets’ amounts ambiguity is not as stark, and this same 50:51 disadvantage may cause an individual to rank a_4 over a_3 . Such preferences clearly violate the Sure-Thing Principle as well as Machina and Schmeidler’s (1992) Strong Comparative Probability Axiom, and are accordingly not compatible with probabilistically sophisticated beliefs. But in addition, they are seen to violate the Comonotonicity Axiom of the Choquet Expected Utility model.

	50 balls		51 balls	
	red	black	green	yellow
a_1	\$45	\$45	\$40	\$40
a_2	\$45	\$40	\$45	\$40
a_3	\$1,000	\$45	\$40	\$0
a_4	\$1,000	\$40	\$45	\$0

⁶⁹This is based on an example of Machina (2009). Baillon et al. (2011) demonstrate that similar examples can also violate standard specifications of the Multiple-Priors, Smooth Ambiguity Preferences (with concave ϕ) and Variational Preferences models, though not the Vector Expected Utility model. (See also Siniscalchi (2008b) as well as the experimental results of L’Haridon and Placido (2010).) Lehrer (2007) demonstrates that they do not violate the models of Lehrer (2009,2012) and Dillenberger and Segal (2012) demonstrate that they do not violate the model of Segal (1987a).

Three particular aspects of ambiguity and ambiguity aversion can arise once three or more outcomes are allowed, each of which poses problems for at least some of the models reviewed in Section 13.5. The first is that some of the models' preferences over *purely subjective* (i.e., not mixed) prospects can inherit the same Allais-type difficulties⁷⁰ faced by expected utility preferences over purely objective prospects. The reason for this is that some purely subjective prospects can be said to be “more objective” than others, and some of the models' preferences over such “almost-objective” prospects converge to expected utility preferences.

To construct such prospects, partition a continuum state space $\mathcal{S} = [0, 1)$ into m equal intervals $\{[0, 1/m), \dots, [i/m, (i+1)/m), \dots, [(m-1)/m, 1)\}$, and for each $\alpha \in [0, 1]$ define $[0, \alpha] \times \mathcal{S}$ as the union of the left α portions of these intervals, so that $[0, \alpha] \times \mathcal{S} \sim_{i=0}^{m-1} [i/m, (i + \alpha)/m)$. As shown by Poincaré (1912) and others,⁷¹ such events will satisfy $\lim_{m \rightarrow \infty} \pi([0, \alpha] \times \mathcal{S}) = \alpha$ for any measure $\pi(\cdot)$ over $[0, 1)$ with a sufficiently regular density. More generally, for any set $\wp \subseteq [0, 1)$ consisting of a finite union of intervals, and any positive integer m , define event $\wp \times \mathcal{S}$ by

$$\wp \times \mathcal{S} = \bigcup_{i=0}^{m-1} \{ (i + \omega)/m \mid \omega \in \wp \}$$

that is, as the union of the natural images of \wp into each of \mathcal{S} 's equal-length intervals. Events with this type of periodic structure are termed *almost-objective events*, and satisfy $\lim_{m \rightarrow \infty} \pi(\wp \times \mathcal{S}) = \lambda(\wp)$ where $\lambda(\cdot)$ is the uniform Lebesgue measure over $[0, 1)$. Bets of the form $(x_1 \text{ on } \wp_1 \times \mathcal{S}; \dots; x_n \text{ on } \wp_n \times \mathcal{S})$ are termed *almost-objective bets*, and are seen to be purely subjective.

As shown in Machina (2013), under natural smoothness conditions, Multiple-Priors,⁷² Smooth Ambiguity Preferences and Variational Preferences preferences over almost-objective bets will converge to expected utility, that is

$$\lim_{m \rightarrow \infty} W_{MP} (x_1 \text{ on } \wp_1 \times \mathcal{S}; \dots; x_n \text{ on } \wp_n \times \mathcal{S}) = \sum_{i=1}^n U(x_i) \cdot \lambda(\wp_i)$$

$$\lim_{m \rightarrow \infty} W_{SM} (x_1 \text{ on } \wp_1 \times \mathcal{S}; \dots; x_n \text{ on } \wp_n \times \mathcal{S}) = \phi (\sum_{i=1}^n U(x_i) \cdot \lambda(\wp_i))$$

$$\lim_{m \rightarrow \infty} W_{VP} (x_1 \text{ on } \wp_1 \times \mathcal{S}; \dots; x_n \text{ on } \wp_n \times \mathcal{S}) = \sum_{i=1}^n U(x_i) \cdot \lambda(\wp_i) + \min_{\pi(\cdot) \in P} c(\pi(\cdot)).$$

⁷⁰ Such as the classic Allais Paradox, Common Consequence Effect and Common Ratio Effect (e.g., Allais (1953), MacCrimmon and Larsson (1968)).

⁷¹ See Machina (2004, p.9).

⁷² For the Multiple-Priors result, the set \mathcal{C} is also assumed to be a convex polytope (the convex hull of a finite number of probability distributions).

Thus, if an individual is asymptotically indifferent between the almost-objective bets

$$\begin{aligned}
 a_1 &= (\$1\text{M}, 1) \\
 a_2 &= (\$5\text{M on } [0, .10] \times S; \$1\text{M on } [.10, .99] \times S; \$0 \text{ on } [.99, 1] \times S) \\
 a_3 &= (\$5\text{M on } [0, .10] \times S; \$0 \text{ on } [.11, 1] \times S) \\
 a_4 &= (\$1\text{M on } [0, .11] \times S; \$0 \text{ on } [.11, 1] \times S)
 \end{aligned}$$

and their respective purely objective counterparts

$$\begin{aligned}
 \hat{a}_1 &= (\$1\text{M}, 1) & \hat{a}_2 &= (\$5\text{M}, .10; \$1\text{M}, .89; \$; 0, .01) \\
 \hat{a}_3 &= (\$5\text{M}, .10; \$0, .90) & \hat{a}_4 &= (\$1\text{M}, .11; \$0, .89)
 \end{aligned}$$

the standard Allais-type preferences $\hat{a}_1 \succ \hat{a}_2$ and $\hat{a}_3 \succ \hat{a}_4$ over these objective bets induce the corresponding preferences $a_1 \succ a_2$ and $a_3 \succ a_4$ over these purely subjective acts, violating Multiple-Priors Preferences, Smooth Ambiguity Preferences and Variational Preferences over purely subjective uncertainty.

A second aspect concerns the Choquet Expected Utility model’s ability to capture attitudes toward different sources, with different amounts, of ambiguity. Take an urn with a single ball, which is either black or white, as well as a coin which has been slightly (but only slightly) bent, and consider the following bets, where the coin is flipped and the ball is drawn simultaneously. Bet I spreads the uncertainty of receiving +\$8,000 versus -\$8,000 across the only slightly ambiguous coin whereas Bet II spreads this uncertainty across the more ambiguous urn, so that ambiguity averter would presumably prefer Bet I.

	Bet I			Bet II	
	black	white		black	white
heads	+\$8,000	\$0	vs.	\$0	\$0
tails	-\$8,000	\$0		-\$8,000	+\$8,000

To see that the Choquet model is incapable of exhibiting such a preference, observe that the model evaluates the bets according to the respective formulas⁷³


$$\begin{aligned}
 \text{Bet I} &: U(\$8,000) \cdot C(HB) + U(-\$8,000) \cdot [1 - C((HB) \cup C(HW) \cup C(TW))] \\
 \text{Bet II} &: U(\$8,000) \cdot C(TW) + U(-\$8,000) \cdot [1 - C((HB) \cup C(HW) \cup C(TW))]
 \end{aligned}$$

On the assumption that the capacity $C(\cdot)$ evaluates the informationally symmetric events HB and TW equally, these values are equal, so the Choquet model must rank the

⁷³ Where HB denotes a heads and a black ball, etc., and we set $U(0) = 0$.

two bets as indifferent, and cannot allow a strict preference over these different sources and amounts of ambiguity.⁷⁴

A final aspect which arises once three or more outcomes are allowed concerns how ambiguity attitudes are allowed to vary with wealth. Defining $\$c$ as the individual's certainty equivalent of an objective 50:50 lottery over $\$0:\100 , the bets on Urns I and II below are each seen to differ from the purely objective bet b_0 by the introduction of ambiguity across some pair of its events. Urn I is obtained from bet b_0 by introducing ambiguity across its middle and lower outcome, whereas Urn II is obtained from b_0 by introducing ambiguity across its middle and upper outcome.



For the same reason that individuals might be less risk averse over higher than lower outcomes, they may exhibit less ambiguity aversion over higher than lower outcomes, which would lead to a preference for Urn II. However, the Multiple-Priors, Rank-Dependent, Smooth and Variational Preferences models evaluate the acts via the respective formulas

$$W_{RD}(\text{URN I}) = 2/3 \cdot C(WW) + 1/2 \cdot [C(WW \cup BW \cup WB) - C(WW)] + 1/3 \cdot [1 - C(WW \cup BW \cup WB)] = W_{RD}(\text{URN II})$$

$$W_{MP}(\text{URN I}) = \min_{(p_{BB}, p_{BW}, p_{WB}, p_{WW}) \in P_0} [1/3 \cdot p_{BB} + 1/2 \cdot p_{BW} + 1/2 \cdot p_{WB} + 2/3 \cdot p_{WW}] = W_{MP}(\text{URN II})$$

$$W_{SM}(\text{URN I}) = \int \phi(1/3 \cdot p_{BB} + 1/2 \cdot p_{BW} + 1/2 \cdot p_{WB} + 2/3 \cdot p_{WW}) \cdot d\mu(p_{BB}, p_{BW}, p_{WB}, p_{WW}) = W_{SM}(\text{URN II})$$

$$W_{VP}(\text{URN I}) = \min_{(p_{BB}, p_{BW}, p_{WB}, p_{WW}) \in P} [1/3 \cdot p_{BB} + 1/2 \cdot p_{BW} + 1/2 \cdot p_{WB} + 2/3 \cdot p_{WW} + c(p_{BB}, p_{BW}, p_{WB}, p_{WW})] = W_{VP}(\text{URN II})$$

⁷⁴ Although not related to the present example, attitudes toward different sources of uncertainty have been experimentally examined by Abdellaoui et al. (2011).

Accordingly, none of the four models can allow attitudes toward ambiguity to depend upon the wealth level at which it occurs.

The reason for this incompatibility comes out when the two bets are expressed in the Anscombe–Aumann format of Section 13.3.4, that is, as mappings from each urn’s four possible states $\{BB, BW, WB, WW\}$ ⁷⁵ to their implied objective lotteries. Setting $U(\$100) = 1$ and $U(\$0) = 0$ so that $U(\$c) = 1/2$, the statewise expected utilities of the two urns are seen to be identical. Since the four models evaluate Anscombe–Aumann acts via their statewise expected utilities, they cannot discriminate between the two urns’ bets.

	<i>BB</i>	<i>BW</i>	<i>WB</i>	<i>WW</i>
URN I	(\$0, 2/3; \$100, 1/3)	(\$0, 1/3; \$c, 1/3; \$100, 1/3)	(\$0, 1/3; \$c, 1/3; \$100, 1/3)	(\$c, 2/3; \$100, 1/3)
URN II	(\$0, 1/3; \$c, 2/3)	(\$0, 1/3; \$c, 1/3; \$100, 1/3)	(\$0, 1/3; \$c, 1/3; \$100, 1/3)	(\$0, 1/3; \$100, 2/3)
expected utility	1/3	1/2	1/2	2/3

A similar example⁷⁶ establishes that the four models cannot express a preference or aversion to ambiguity in losses versus gains:

Ambiguity in Losses versus Gains									
URN I				URN II					
1 ball		1 ball		1 ball		1 ball			
black	white	red	green	vs.	red	green	black	white	
-3	-1	+2	+2		-2	-2	+1	+3	

Another recent variation on Ellsberg’s standard form was examined by Yang and Yao (2011), who allowed subjects to draw twice with replacement from the same urn, with a different color winning each time, again finding a substantial number of violations of the major theories of ambiguity aversion.

⁷⁵ Where *BW* denotes that the first ball in the urn is black and the second is white, etc.

⁷⁶ With outcomes already expressed in utils.

13.7 UPDATING AND DYNAMIC CHOICE

Savage's axiomatization of expected utility is atemporal—time is not an explicit consideration or an input to the theory. Alternatively, one can view Savage's theory as dealing with one-period decision problems: the individual chooses an act, then the state is realized, which determines the prize accruing to the decision maker. However, Postulate P2, the Sure-Thing Principle, provides a way to extend the theory to *dynamic decision problems*, in which the individual may acquire partial information about the state of nature over time, and take actions at several decision points, until all uncertainty is resolved and a final prize is obtained. A formal analysis of the extension of Savage's atemporal theory to dynamic choice is beyond the scope of this chapter; instead, we emphasize two key ingredients of this extension.

The first is *updating*. The notion of updating a prior probability measure μ so as to reflect information the individual has acquired is familiar. If one assumes that the individual is an expected utility maximizer, and that his or her risk attitudes and tastes do not change, the probabilistic conditioning operator determines a *conditional preference*. Savage (1954, p.22) suggested the converse approach: one can start with a *definition* of conditional preference, and show that it characterizes probabilistic updating. Specifically, Savage stipulates that act f is weakly preferred to act g given event E , written $f \succcurlyeq_E g$, if and only if, for *some* act h ,

$$(f, E; h, \sim E) \succcurlyeq (g, E; h, \sim E). \quad (13.20)$$

It is easy to see that, if the *ex-ante* preference \succcurlyeq is an *EU* preference, with subjective prior μ , then the updated preference \succcurlyeq_E is also consistent with *EU*. Then, if the event E is nonnull,⁷⁷ one has $\mu(E) > 0$, and the probability measure associated with the updated preference \succcurlyeq_E is the usual update $\mu(\cdot|E)$ of μ . We emphasize that Postulate P2 is essential to ensure that this definition is well-posed: if P2 is violated, then different choices of acts h may lead to contradictory rankings of f and g conditional on E .

The second key ingredient in the extension of Savage's theory to dynamic choice is *dynamic consistency*. One possible formulation of this property is as follows: if act f is weakly preferred to act g conditional on E , and also conditional on $\sim E$, then f is unconditionally preferred to g ; also, if one of the conditional preferences is strict, so is the unconditional preference. Formally:

$$f \succcurlyeq_E g, f \succcurlyeq_{\sim E} g \Rightarrow f \succcurlyeq g \quad \text{and} \quad f \succ_E g, f \succcurlyeq_{\sim E} g \Rightarrow f \succ g \quad (13.21)$$

(The second part of the definition requires that the event E be nonnull.) Dynamic consistency is essential to ensure that, if a multiperiod choice problem is solved by *backward*

⁷⁷ See footnote 25.

induction or *recursion*, one obtains a solution that is also optimal *ex ante*. It turns out that Savage’s Postulate P2 also ensures that conditional preferences as defined above satisfy dynamic consistency. (Indeed, the connection between P2 and dynamic consistency is tight: see Epstein and Le Breton (1993), Ghirardato (2002) and Siniscalchi (2011, Sect. 4.2).)

Since ambiguity-sensitive preferences violate Postulate P2, it should come as no surprise that updating and consistent dynamic choice pose challenges. As we shall see, there is no unique way to define conditional preferences or—in any given parametric model—conditional “beliefs.” Furthermore, no matter how we define conditional preferences, we can construct examples in which dynamic consistency is violated. Different approaches have been proposed to deal with these issues. We now briefly discuss some of the relevant contributions.

Lack of space prevents us from discussing two other broadly related issues which have seen recent work. The first is learning under ambiguity: see Epstein and Schneider (2007). The second is the notion of exchangeability, or symmetry: Epstein and Seo (2010, 2012) extend the classic results of de Finetti (1937) and Hewitt and Savage (1955) to ambiguity-sensitive preferences. See also Al Najjar and De Castro (2013).

13.7.1 Updating Ambiguous Beliefs

Early contributions focus on updating rules for capacities in the Choquet Expected Utility model. Two rules have received particular attention: the *Dempster-Shafer Rule* and the *Full Bayesian Rule*. Both essentially originate in the work of Dempster (1967), who considers the following scenario: An individual has a probabilistic belief μ on a set Ω ; furthermore, there is a correspondence (a multivalued map) $\Gamma : \Omega \rightarrow \mathcal{S}$, where \mathcal{S} is the state space of interest, which for simplicity we assume to be finite. Intuitively, the map Γ describes the individual’s partial information about the relationship between elements of Ω , whose likelihood the individual can assess with confidence, and the events in \mathcal{S} . The individual’s beliefs on \mathcal{S} are represented by the capacity ν defined by

$$\nu(A) = \frac{\mu(\{\omega : \emptyset \neq \Gamma(\omega) \subseteq A\})}{\mu(\{\omega : \Gamma(\omega) \neq \emptyset\})}.$$

Intuitively, this embodies a conservative stance: the individual assigns to each A only the probability of those elements of Ω that he or she knows for sure to correspond to elements in A . The capacity thus defined is convex, and indeed satisfies a stronger property that characterizes so-called “belief functions” (e.g., Shafer (1976) and Jaffray (1992)). The reason why Γ is allowed to take on an empty value will be clear momentarily.

Dempster describes the following updating rule. Suppose the individual learns that event $E \subseteq \mathcal{S}$ has occurred. The initial information, as encoded by Γ , was that element ω of Ω applied to states s in $\Gamma(\omega)$: but since he or she now knows that the state s must also

lie in E , he or she associates ω with $\Gamma(\omega) \cap E$. We can then define $\Gamma_E(\omega) = \Gamma(\omega) \cap E$ and construct the capacity $\nu_{DS}(\cdot|E)$ derived from Γ_E as above (we explain the superscript below). Dempster shows that

$$\nu_{DS}(A|E) = \frac{\nu([A \cap E] \cup [\sim E]) - \nu(\sim E)}{1 - \nu(\sim E)}.$$

While this expression may seem a bit mysterious, the above construction clarifies its origin. The subsequent literature has mostly taken the above formula (or alternative, equivalent formulae) as the *definition* of a particular updating rule; since it also plays a prominent role in the work of Shafer (1976), it is commonly referred to as the *Dempster-Shafer* updating rule (hence the subscript in $\nu_{DS}(\cdot|E)$).

An alternative updating rule for capacities is related to sets of priors. It is known that, if a capacity ν is a “belief function,” then it satisfies the following duality relation: letting \mathcal{C}_ν be the set of probability distributions μ such that $\mu(A) \geq \nu(A)$ for all events A , it is the case that, for every event A , $\nu(A) = \inf\{\mu(A) : \mu \in \mathcal{C}_\nu\}$ (this is *not* true for arbitrary convex capacities). Thus, a possible updating rule is as follows: first, assuming $\nu(E) > 0$, update all probability distributions in \mathcal{C}_ν according to the usual probabilistic conditioning formula; then, consider the lower envelope of the resulting set. That is, we define the so-called *Full Bayesian updating rule* for the belief function ν by $\nu_{FB}(A|E) = \inf\{\mu(A|E) : \mu \in \mathcal{C}_\nu\}$. Jaffray (1992) shows that $\nu_{FB}(\cdot|E)$ is indeed a belief function, and furthermore that

$$\nu_{FB}(A|E) = \frac{\nu(A)}{\nu(A) + 1 - \nu(A \cup \sim E)}.$$

An alternative approach leading to the above definition of Full Bayesian updating can be found in Fagin and Halpern (1991).

Gilboa and Schmeidler (1993) are the first to provide an axiomatic treatment of updating rules in a decision setting; in particular, they adopt the horse-roulette act environment discussed in Section 13.5, and assume that there exist a best and a worst prize. They consider uncertainty-averse Choquet Expected Utility preferences (which are thus also Maxmin Expected Utility preferences) and consider the class of *h-Bayesian update rules*, where h is an arbitrary act. Given an unconditional preference \succsim , the *h-Bayesian update* given E , denoted $\succsim_{E,h}$, is defined by

$$f \succsim_{E,h} g \Leftrightarrow (f, E; h, \sim E) \succsim (g, E; h, \sim E).$$

Under Savage’s Postulate P2, all *h-Bayesian update rules* coincide with Savage’s update rule; however, for more general preferences, different choices of h induce different conditional preferences. In particular, Gilboa and Schmeidler show that, if one takes

h to be the constant act that yields the *best possible prize*, the resulting conditional preference is *CEU*, with capacity given by the Dempster-Shafer update of the prior capacity (the utility function is unchanged relative to the prior preference). They interpret this as a form of “pessimistic updating”: when receiving the information that E has occurred, the individual ranks acts as if, in the counterfactual event, he or she would have received the best possible prize.

Eichberger et al. (2007) and Horie (2007) characterize Full Bayesian updating for arbitrary capacities. Their main axiom is adopted from Pires (2002), a paper we discuss below; thus, we defer the discussion of their contribution.

Two main updating rules have been proposed for maxmin expected utility (*MEU*) preferences. The first stipulates that, given a set \mathcal{C} of priors on the state space \mathcal{S} , all priors $\mu \in \mathcal{C}$ with $\mu(\cdot|E) > 0$ be updated upon learning that E has occurred; this is called the *Full* or *Generalized Bayesian updating rule* for the set \mathcal{C} . It is related to Full Bayesian updating for belief functions, in the sense that (i) a *CEU* preference for which the capacity is a belief function, and hence convex, is also a *MEU* preference, and (ii) by the result stated above, the *CEU* preference characterized by the full Bayesian update of the capacity is the same as the *MEU* preference characterized by the full Bayesian update of the set \mathcal{C} . However, it is important to emphasize that full Bayesian updating of sets of priors is well-defined for arbitrary sets \mathcal{C} , not just those that consist of probability distributions that dominate a given belief function.

Full Bayesian updating of priors has been axiomatized by Pires (2002) in the horse-roulette act setup; related results can be found in Walley (1991, pp. 632–638) and Jaffray (1994). The key axiom establishes a connection between the conditional certainty equivalent x of an act f and unconditional ranking of x *vis-à-vis* a suitable composite act: formally, for every (nonnull) event E , act f and prize x , it requires that $f \sim_{E,x}$ if and only if $(f, E; x, \sim E) \sim x$. One implication of this axiom is that conditional preferences can be elicited from prior preferences: this is done by determining, for each act f , the prize x that solves the preference equation $(f, E; x, \sim E) \sim x$.⁷⁸

Eichberger et al. (2007) use the same axiom to characterize Full Bayesian Updating for capacities; however, Horie (2007) points out that, while Pires’s axiom is sufficient to deliver full Bayesian updating, it is too strong. Horie shows that the appropriate necessary and sufficient condition is a version of Pires’s axiom that is restricted to binary acts f .

An alternative updating rule was proposed and characterized by Gilboa and Schmeidler (1993). Recall that they characterize Dempster-Shafer updating for uncertainty averse *CEU* preferences, which are thus also *MEU* preferences. Hence, their result necessarily also pins down a procedure to update sets of priors. The procedure


⁷⁸ Siniscalchi (2011) shows that Pires’s axiom is equivalent to a weakening of the standard dynamic consistency property discussed above.

works as follows: out of the set \mathcal{C} of priors characterizing the *ex-ante* preference, select the measures assigning the maximum probability to the event E that has occurred, and only update those. Gilboa and Schmeidler call this the *maximum-likelihood update rule*. Note that their characterization result only covers preferences that are both *MEU* and *CEU*; a characterization of maximum-likelihood updating for general *MEU* preferences is not known.

Finally, Siniscalchi (2009a) defines and axiomatizes an updating rule for *VEU* preferences, which entails (i) updating the baseline prior in the usual way, and (ii) replacing the covariances in equation (13.10) with suitably rescaled conditional covariances.

13.7.2 Dynamic Choice under Ambiguity

Regardless of the preference model and updating rule one adopts, the potential for dynamic inconsistency of preferences arises as soon as one relaxes Savage's Postulate P2. This is easy to see from the following decision tree, which can be thought of as a dynamic version of the Three-Color Paradox. There are two decision points, identified by filled circles; in the first, *ex-ante*, stage, the individual chooses whether to stop (S) and get a 50% chance to receive \$100, or continue (C) with the bet. Then, the individual learns whether the ball drawn is yellow or not; if it is yellow, the individual receives $\$x$. If instead the ball is not yellow, the individual can choose to bet on red or black. The composition of the urn, and the individual's information about it, is as in the (static) Three-Color Paradox.



A Dynamic Three-Color Urn Paradox

The objects of choice in this problem are not, strictly speaking, acts: rather, they are *actions* (such as R and B , C and S) at different nodes. The individual can also formulate *plans*, such as “choose C , then choose R if the second decision node is reached” (denoted CR). The atemporal models discussed in Section 13.5 do not directly induce preferences over actions or plans. However, one can adopt a standard *reduction* assumption, and evaluate a plan or action according to the act that it induces. For instance, the plan CR induces the act $(\$100, s_r; \$0, s_b; \$x, s_y)$; this can be evaluated by any one of the

functional representations of Section 13.5. Similarly, the action R at the second decision node induces a “partial” act $(\$100, s_r; \$0, s_b)$; any one of the updating rules discussed in the previous subsection can be used to evaluate such acts, as prizes assigned on the event s_y are irrelevant when conditioning on $\{s_r, s_b\}$.

To see that dynamic consistency may fail, suppose *ex-ante* preferences are *MEU*, with priors as in equation (13.2). Consider the case $x = \$100$, and observe that then the plans CR and CB correspond to the Three-Color urn’s bets a_3 and a_4 respectively. As noted in Section 13.5.1, letting $U(\$100) = 1$ and $U(\$0) = 0$, bet a_4 , and therefore CB yield utility $2/3$, whereas a_3 and hence CR yield utility $1/3$; on the other hand, choosing S yields utility $1/2$. Therefore, the *ex-ante* optimal plan for this decision maker is CB . Now suppose that this individual uses the Full Bayesian updating rule. The updated set of probability distributions is

$$C_{\{s_r, s_b\}} = \{ \mu : \mu(\{s_r\}) \geq 1/3, \mu(\{s_y\}) = 0 \}.$$

Therefore, upon reaching the second decision node, the individual will strictly prefer R , which yields $\$100$ with probability at least $1/3$, to B , which may yield $\$100$ with probability 0 . This means that, *although CR is the ex-ante optimal plan, the individual will not actually want to carry it out.* Furthermore, if we analyze the problem according to backward induction, we conclude that the individual will choose R at the second node. Hence, from the perspective of the first node, choosing C is “just as good” as committing to plan CB , which—as noted above—has an *ex-ante* utility of only $1/3$. Therefore, the individual will prefer to choose S at the initial node, according to backward induction. Thus, *backward induction yields a different solution than ex-ante optimization.* Both these issues can be traced back to a violation of dynamic consistency: the acts a_4 and a_3 (equivalently, the plans CR and CB) are indifferent conditional on $\{s_y\}$ because they both yield $\$x = \100 , the individual strictly prefers a_3 to a_4 conditional on $\{s_r, s_b\}$, but strictly prefers a_4 to a_3 unconditionally.

Similar examples can be constructed for different preference representations and updating rules; the issues highlighted are not unique to *MEU* or full Bayesian updating.

There are three ways to address this failure of dynamic consistency.⁷⁹ The first was introduced by Epstein and Schneider (2003) in the context of *MEU* preferences, and extended to Variational Preferences by Maccheroni et al. (2006b), to Smooth Ambiguity Preferences by Klibanoff et al. (2009), and to Vector Expected Utility preferences by Siniscalchi (2010). This approach has also been broadly adopted in applications, especially in financial economics. The basic observation is that the noted tight connection

⁷⁹ A fourth way is to drop the reduction assumption. In the example, one could argue that the individual may prefer plan CR to plan CB in the tree of the figure, even though he or she prefers act a_4 to act a_3 in the static Three-Color Paradox. This is explored in Li (2011).

between Savage's Postulate P2 and dynamic consistency does *not* prevent the coexistence of *some* departures from P2 and *some* form of dynamic consistency. More precisely, if for a given state space one wishes to ensure dynamic consistency in *every possible decision tree*, then Postulate P2 must hold for *all* events, and therefore there is no scope for Ellsberg behavior. However, one can fix *one* particular decision tree, and require dynamic consistency only for that particular tree. This implies that P2 will hold only for *some* events, but there will still be scope for Ellsberg-type behavior.

For a simple example, suppose that a fair coin is tossed; if it comes up heads, the individual is confronted with the Three-Color Ellsberg urn and must choose to bet on red vs. black; if it comes up tails, the individual faces the same urn but must choose whether to bet on red or yellow vs. black or yellow. Letting the space be $\mathcal{S} = \{H, T\} \times \{r, b, y\}$ to simplify notation, a plausible set of priors for this problem is

$$\mathcal{C} = \left\{ \mu : \mu(\{H_r\}) = \mu(\{T_r\}) = 1/6, \mu(\{H_b\}) = \mu(\{T_b\}), \mu(\{H_y\}) = \mu(\{T_y\}) \right\}.$$

For this set of priors, the modal Ellsberg preferences obtain, both *ex ante* and after each realization of the coin toss. Dynamic consistency also holds in this tree. Note that this set of priors is obtained by taking the product of the uniform probability on $\{H, T\}$ with the priors in the set \mathcal{C} of equation (13.2). Epstein and Schneider show that dynamic consistency of *MEU* preferences in a particular tree characterizes a generalization of this property of the representing set of priors, called *rectangularity*. Furthermore, they show that dynamic *MEU* preferences with a rectangular set of priors admit a *recursive* representation; this accounts for the popularity of this approach in applications.

Rectangularity does rule out, for example, the set of priors in equation (13.2) from Section 13.5.1, which we used to analyze the dynamic version of the Three-Color Paradox in the above figure. Intuitively, conditioning on the event $\{s_r, s_b\}$ “breaks” the *complementarity* between the states s_b and s_y (cf. Sections 13.5.6 and 13.6.1) and this leads to a violation of dynamic consistency. Loosely speaking, rectangularity rules out precisely these complementarities across different conditioning events. Of course, in some contexts, such complementarities may be the main object of interest—as is the case in the stylized problem of the figure. In such cases, alternatives to the Epstein-Schneider approach can be pursued.

One such alternative was advocated by Hanany and Klibanoff (2007, 2009), who adapt arguments first proposed by Machina (1989) in the context of risky choice. The basic intuition is that the preceding arguments implicitly made the assumption that, when the individual is evaluating choices at a node in a decision tree, *bygones are bygones*: the only relevant considerations are the residual uncertainty, and the prizes delivered at terminal nodes that can still be reached. Machina, as well as Hanany and Klibanoff, argues that this need not be the case. For example, in the choice problem of the figure, when the individual faces a choice between *R* and *B*, he or she may take into account

that this choice occurs “in the context” of a given decision tree. In particular, he or she may take into account *the prize* $\$x$ *that would have been delivered if a yellow ball had been drawn*. Formally, this allows the individual to express a *conditional* preference for R over B if $x = \$0$, and for B over R if $x = \$100$. Hence, in particular, with $x = \$100$ there no longer is any dynamic inconsistency: *ex ante*, the individual prefers CB to CR , and upon reaching the second decision node, prefers B to R .

Formally, Hanany and Klibanoff (2007) focus on *MEU* preferences, and impose a (weak) form of dynamic consistency. They consider a class of two-period decision problems and a class of updating rules, i.e., mappings that associate a second-period conditional preference to every prior preference, decision problem and *ex-ante* optimal act. Within this class, they characterize the ones that are dynamically consistent in the sense that the *ex-ante* optimal act remains conditionally optimal. Their 2009 paper extends the results to arbitrary Uncertainty-Averse Preferences.

One consequence of this approach is that the resulting conditional preferences may display either uncertainty aversion or uncertainty neutrality, depending on the prizes assigned at counterfactual terminal nodes; an example is provided in Siniscalchi (2009b). This seems at odds with the usual interpretation of ambiguity, which concerns information rather than prizes.

The final approach we discuss can be traced back to Strotz (1956), and was advocated in the context of ambiguity by Siniscalchi (2011). The basic idea is to assume that individuals are *sophisticated*: that is, even though they may be dynamically inconsistent, they correctly anticipate their future behavior, and respond accordingly. Strotz argues that a sophisticated individual should engage in a *strategy of consistent planning*: at each decision point, they should choose the plan that is best according to their *current* preferences, among all those plans that, in view of their future preferences, they will *actually* be able to carry out. Consistent planning is thus a “refinement” of backward induction, in the sense that, whenever two or more plans are optimal from the perspective of future preferences, the one that is optimal for the current preferences is selected.

In the tree of the figure with $x = \$100$, consistent planning implies that, at the initial node, the individual will anticipate their preference for R over B upon learning that a yellow ball was not drawn; hence, they will realize that choosing C is tantamount to committing to the plan CR , and will therefore choose S instead. In other words, consistent planning coincides with backward induction in this example.

Siniscalchi provides an axiomatization of consistent planning. The key insight is that, to formulate the sophistication principle (and other assumptions implicit in consistent planning), preferences need to be defined over *decision trees*, rather than acts, or even plans. This is analogous to the use of preferences over *menus* in the literature on

preference for flexibility, or temptation and self-control. One advantage of this approach is that it yields insight into the issue of the *value of information* under ambiguity. The reader is referred to the original paper for detail.

13.8 CONCLUSION

In recent decades the theoretical and experimental literature on the economics of ambiguity and ambiguity aversion spawned by Ellsberg (1961, 1962) has exploded to the point where it is much too extensive to be covered in a single chapter of this form. For additional discussion, the reader is referred to the surveys of Sugden (1986), Camerer and Weber (1992), Kelsey and Quiggin (1992), Kischka and Puppe (1992), Camerer (1995), Mukerji (2000), Starmer (2000), Siniscalchi (2008a), Gilboa et al. (2008), Al-Najjar and Weinstein (2009), Hey et al. (2010), Etner et al. (2012) and Gilboa and Marinacci (2012), the journal symposium issues of Bonanno et al. (2009) and Ellsberg et al. (2011), and the extensive annotated bibliography of Wakker (2013).

REFERENCES

- Abdellaoui, M., Baillon, A., Placido, L., Wakker, P., 2011. The rich domain of uncertainty: source functions and their experimental implementation. *American Economic Review* 101, 695–723.
- Ahn, D.S., 2008. Ambiguity without a state space. *Review of Economic Studies* 75, 3–28.
- Al Najjar, N., De Castro, L., 2013. Parametric Representation of Preferences. manuscript., Kellogg Graduate School of Management, Northwestern University.
- Al-Najjar, N., Weinstein, J., 2009. The ambiguity aversion literature: a critical assessment. *Economics and Philosophy* 25, 249–284.
- Allais, M., 1953. Le comportement de l'homme rationnel devant le risque, critique des postulats et axiomes de l'école Américaine. *Econometrica* 21, 503–546.
- Anscombe, F., Aumann, R., 1963. A definition of subjective probability. *Annals of Mathematical Statistics* 34, 199–205.
- Arrow, K., 1951. Alternative approaches to the theory of choice in risk-taking situations. *Econometrica* 19, 404–437.
- Arrow, K., 1963. Comment. *Review of Economics and Statistics* 45 (Supplement), 24–27.
- Baillon, A., Bleichrodt, H., 2011. Testing Ambiguity Models Through the Measurement of Probabilities for Gains and Losses. manuscript. Erasmus University Rotterdam.
- Baillon, A., Cabantous, L., Wakker, P., 2012. Aggregating imprecise or conflicting beliefs: an experimental investigation using modern ambiguity theories. *Journal of Risk and Uncertainty* 44, 115–147.
- Baillon, A., l'Haridon, O., Placido, L., 2011. Ambiguity models and the machina paradoxes. *American Economic Review* 101, 1547–1560.
- Bayes, T., 1763. An essay toward solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London* 53, 370–418.
- Becker, S., Brownson, F., 1964. What price ambiguity? or the role of ambiguity in decision-making. *Journal of Political Economy* 72, 62–73.
- Becker, G., DeGroot, M., Marschak, J., 1964. Measuring utility by a single-response sequential method. *Behavioral Science* 9, 226–232.
- Bernoulli, D., 1738. *Specimen Theoriae Novae De Mensura Sortis, Commentarii Academiae Scientiarum Imperialis Petropolitanae* (Papers of the Imperial Academy of Sciences in Petersburg), 175–192. English Translation: Exposition of a New Theory on the Measurement of Risk. *Econometrica* 22 (1954), 23–36.
- Bewley, T., 2002. Knightian decision theory. Part 1. *Decisions in Economics and Finance* 25, 79–110.
- Bewley, T., 2011. Market innovation and entrepreneurship: a Knightian view. In: Debreu, G., Neufeuind, W., Trockel, W. (Eds.), *Economics Essays: A Festschrift for Werner Hildenbrand*. Springer.
- Bonanno, G., Van Hees, M., List, C., Tungodden, B., 2009. Special issue on ambiguity aversion. *Economics and Philosophy* 25, 247–369.
- Brewer, K., 1963. Decisions under uncertainty: comment. *Quarterly Journal of Economics* 77, 159–161.

- Brewer, K., Fellner, W., 1965. The slanting of subjective probabilities—agreement on some essentials. *Quarterly Journal of Economics* 79, 657–663.
- Camerer, C., 1995. Individual decision making. In: Kagel, J., Roth, A. (Eds.), *Handbook of Experimental Economics*. University Press, Princeton.
- Camerer, C., Kunreuther, H., 1989a. Decision processes for low probability events: policy implications. *Journal of Policy Analysis and Management* 8, 565–592.
- Camerer, C., Kunreuther, H., 1989b. Experimental markets for insurance. *Journal of Risk and Uncertainty* 2, 265–300.
- Camerer, C., Weber, M., 1992. Recent developments in modeling preferences: uncertainty and ambiguity. *Journal of Risk and Uncertainty* 5, 325–370.
- Carter, C., 1950. Expectation in Economics (Review of Shackle (1949a)). *Economic Journal* 60, 92–105.
- Casadesus-Masanell, R., Klibanoff, P., Ozdenoren, E., 2000a. Maxmin expected utility over savage acts with a set of priors. *Journal of Economic Theory* 92, 33–65.
- Casadesus-Masanell, R., Klibanoff, P., Ozdenoren, E., 2000b. Maxmin expected utility through statewise combinations. *Economics Letters* 66, 49–54.
- Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., Montrucchio, L., 2011a. Uncertainty averse preferences. *Journal of Economic Theory* 146, 1275–1330.
- Cerreia-Vioglio, S., Ghirardato, P., Maccheroni, F., Marinacci, M., Siniscalchi, M., 2011b. Rational preferences under ambiguity. *Economic Theory* 48, 341–375.
- Chambers, R., Grant, S., Polak, B., Quiggin, J., 2013. A Two-Parameter Model of Dispersion Aversion. manuscript, University of Queensland.
- Chateauneuf, A., Faro, J., 2009. Ambiguity through confidence functions. *Journal of Mathematical Economics* 45, 535–558.
- Chernoff, H., 1954. Rational selection of decision functions. *Econometrica* 22, 422–443.
- Chew, S., Li, K., Chark, R., Zhong, S., 2008. Source preference and ambiguity aversion: models and evidence from behavioral and neuroimaging experiments in neuroeconomics. *Advances in Health Economics and Health Services Research* 20, 179–201.
- Chew, S.H., Sagi, J., 2006. Event exchangeability: probabilistic sophistication without continuity or monotonicity. *Econometrica* 74, 771–786.
- Chew, S.H., Sagi, J., 2008. Small worlds: modeling attitudes toward sources of uncertainty. *Journal of Economic Theory* 139, 1–24.
- Chipman, J., 1958. Stochastic choice and subjective probability (abstract). *Econometrica* 26, 613.
- Chipman, J., 1960. Stochastic choice and subjective probability. In: Willner, D. (Ed.), *Decisions Values and Groups*. Pergamon Press, New York.
- Chow, C., Sarin, R., 2001. Comparative ignorance and the Ellsberg paradox. *Journal of Risk and Uncertainty* 22, 129–139.
- Clarke, F.H., 1983. *Optimization and Nonsmooth Analysis*. Wiley-Interscience, New York.
- Cohen, M., Jaffray, J.-Y., Said, T., 1985. Individual behavior under risk and under uncertainty: an experimental study. *Theory and Decision* 18, 203–228.
- Cohen, M., Jaffray, J.-Y., Said, T., 1987. Individual behavior under risk and under uncertainty for gains and losses. *Organizational Behavior and Human Decision Processes* 39, 1–22.
- Curley, S., Eraker, S., Yates, F., 1984. An investigation of patient's reactions to therapeutic uncertainty. *Medical Decision Making* 4, 501–511.
- Curley, S., Yates, F., 1985. The center and range of the probability interval as factors affecting ambiguity preferences. *Organizational Behavior and Human Decision Processes* 36, 272–287.
- Curley, S., Yates, F., 1989. An empirical evaluation of descriptive models of ambiguity reactions in choice situations. *Journal of Mathematical Psychology* 33, 397–427.
- Curley, S., Young, M., Yates, F., 1989. Characterizing physicians' perceptions of ambiguity. *Medical Decision Making* 9, 116–124.
- Davidson, D., Suppes, P., 1956. A finitistic axiomatization of utility and subjective probability. *Econometrica* 24, 264–275.
- de Finetti, B., 1937. La prévision: ses lois logiques, ses sources subjectives. *Annales De L'institut Henri Poincaré* 7, 1–68.
- DeGroot, M., 1970. *Optimal Statistical Decisions*. McGraw-Hill Book Co, New York.

- Dempster, A., 1967. Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics* 38, 325–339.
- Dillenberger, D., Segal, U., 2012. Recursive Ambiguity and Machina's Examples. manuscript, University of Pennsylvania.
- Du, N., Burescu, D., 2005. The effects of imprecise probabilities and outcomes in evaluating investment options. *Management Science* 51, 1791–1803.
- Dumav, M., Stinchcombe, M., 2013. The von Neumann–Morgenstern Approach To Ambiguity. manuscript, University of Texas at Austin.
- Eichberger, J., Grant, S., Kelsey, D., 2007. Updating choquet beliefs. *Journal of Mathematical Economics* 43, 888–899.
- Eichberger, J., Grant, S., Kelsey, D., Koshevoy, G.A., 2011. The meu model: a comment. *Journal of Economic Theory* 146, 1684–1698.
- Einhorn, H., Hogarth, R., 1986. Decision making under ambiguity. *Journal of Business* 59 (Supplement), S225–S250.
- Eliasz, K., Ortoreva, P., 2011. A Variation on Ellsberg. manuscript, Brown University.
- Ellsberg, D., 1961. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics* 75, 643–669.
- Ellsberg, D., 1962. Risk, Ambiguity and Decision. Ph.D. thesis, Harvard University.
- Ellsberg, D., 1963. Risk, ambiguity, and the savage axioms: reply. *Quarterly Journal of Economics* 77, 336–342.
- Ellsberg, D., 2001. Risk, Ambiguity and Decision. Garland Publishing, Inc., New York & London (Published version of Ellsberg (1962)).
- Ellsberg, D., Machina, M., Ritzberger, K., Yannelis, N. (Eds.), 2011. *Economic Theory* 48, 219–548. Symposium on the 50th Anniversary of the Ellsberg Paradox.
- Epstein, L., 1999. A definition of uncertainty aversion. *Review of Economic Studies* 66, 579–608.
- Epstein, L., Le Breton, M., 1993. Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory* 61, 1–22.
- Epstein, L., Schneider, M., 2003. Recursive multiple-priors. *Journal of Economic Theory* 113, 1–31.
- Epstein, L.G., Schneider, M., 2007. Learning under ambiguity. *The Review of Economic Studies* 74, 1275–1303.
- Epstein, L., Seo, K., 2010. Symmetry of evidence without evidence of symmetry. *Theoretical Economics* 5, 313–368.
- Epstein, L., Seo, K., 2012. Ambiguity with Repeated Experiments. manuscript, Boston University.
- Epstein, L., Zhang, J., 2001. Subjective probabilities on subjectively unambiguous events. *Econometrica* 69, 265–306.
- Ergin, H., Gul, F., 2009. A theory of subjective compound lotteries. *Journal of Economic Theory* 144, 899–929.
- Etner, J., Jaleva, M., Tallon, J.-M., 2012. Decision theory under ambiguity. *Journal of Economic Surveys* 26, 234–270.
- Fagin, R., Halpern, J., 1991. A new approach to updating beliefs. In: *Uncertainty in Artificial Intelligence*, 6, pp. 347–374.
- Fellner, W., 1961. Distortion of subjective probabilities as a reaction to uncertainty. *Quarterly Journal of Economics* 75, 670–689.
- Fellner, W., 1963. Slanted subjective probabilities and randomization: reply to Howard Raiffa and K.R.W. Brewer. *Quarterly Journal of Economics* 77, 676–690.
- Fishburn, P., 1970. *Utility Theory for Decision Making*. John Wiley & Sons, New York.
- Fox, C.R., Tversky, A., 1995. Ambiguity aversion and comparative ignorance. *Quarterly Journal of Economics* 110, 585–603.
- Gajdos, T., Hayashi, T., Tallon, J.M., Vergnaud, J.C., 2008. Attitude toward imprecise information. *Journal of Economic Theory* 140, 27–65.
- Georgescu-Roegen, N., 1954. Choice, expectations and measurability. *Quarterly Journal of Economics* 58, 527–530.
- Georgescu-Roegen, N., 1958. The nature of expectation and uncertainty. In: Bowman, M. (Ed.), *Expectations, Uncertainty, and Business Behavior*. Social Science Research Council, New York.
- Gerrity, M., Devellis, R., Earp, J., 1990. Physicians' reactions to uncertainty in patient care: a new measure and new insights. *Medical Care* 28, 724–736.

- Ghirardato, P., 2002. Revisiting savage in a conditional world. *Economic Theory* 20, 83–92.
- Ghirardato, P., Maccheroni, F., Marinacci, M., 2004. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118, 133–173.
- Ghirardato, P., Maccheroni, F., Marinacci, M., Siniscalchi, M., 2003. A subjective spin on roulette wheels. *Econometrica* 71, 1897–1908.
- Ghirardato, P., Marinacci, M., 2002. Ambiguity made precise: a comparative foundation. *Journal of Economic Theory* 102, 251–289.
- Ghirardato, P., Siniscalchi, M., 2010. **A More Robust Definition of Multiple Priors.** Carlo Alberto Notebook N. 144.
- Ghirardato, P., Siniscalchi, M., 2012. Ambiguity in the small and in the large. *Econometrica* 80, 2827–2847.
- Gilboa, I., 1987. Expected utility with purely subjective non-additive probabilities. *Journal of Mathematical Economics* 16, 65–88.
- Gilboa, I., Marinacci, M., 2012. Ambiguity and the Bayesian paradigm. In: *Advances in Economics and Econometrics: Theory and Applications*, Tenth World Congress of the Econometric Society.
- Gilboa, I., Postlewaite, A., Schmeidler, D., 2008. Probability and uncertainty in economic modeling. *Journal of Economic Perspectives* 22, 173–188.
- Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics* 18, 141–153.
- Gilboa, I., Schmeidler, D., 1993. Updating ambiguous beliefs. *Journal of Economic Theory* 59, 33–49.
- Graaf, J., Baumol, W., 1949. Note on expectation in economics. *Economica* 16, 338–342.
- Grant, S., 1995. Subjective probability without monotonicity: or how Machina's mom may also be probabilistically sophisticated. *Econometrica* 63, 159–189.
- Grant, S., Polak, B., 2013. Mean–dispersion preferences and constant absolute uncertainty aversion. *Journal of Economic Theory* 148, 1361–1398.
- Gul, F., 1992. Savage's theorem with a finite number of states. *Journal of Economic Theory* 57, 99–110.
- Gul, F., Pesendorfer, W., 2013. Expected Uncertainty Utility Theory.** manuscript, Princeton University.
- Halevy, Y., 2007. Ellsberg revisited: an experimental study. *Econometrica* 75, 503–536.
- Hanany, E., Klibanoff, P., 2007. Updating preferences with multiple priors. *Theoretical Economics* 2, 261–298.
- Hanany, E., Klibanoff, P., 2009. Updating ambiguity averse preferences. In: *BE Journal of Theoretical Economics* 9.
- Hansen, L., Sargent, T., 2001. Robust control and model uncertainty. *American Economic Review* 91, 60–66.
- Hansen, L., Sargent, T., Tallarini, T., 1999. Robust permanent income and pricing. *Review of Economic Studies* 66, 873–907.
- Hewitt, E., Savage, L., 1955. Symmetric measures on Cartesian products. *Transactions of the American Mathematical Society* 80, 470–501.
- Hey, J., Lotito, G., Maffioletti, A., 2010. The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity. *Journal of Risk and Uncertainty* 41, 81–111.
- Ho, J., Keller, L.R., Keltyka, P., 2002. Effects of outcome and probabilistic ambiguity on managerial choices. *Journal of Risk and Uncertainty* 24, 47–74.
- Hogarth, R., 1989. Ambiguity and competitive decision making: some implications and tests. *Annals of Operations Research* 19, 31–50.
- Hogarth, R., Einhorn, H., 1990. Venture theory: a model of decision weights. *Management Science* 36, 780–803.
- Hogarth, R., Kunreuther, H., 1985. Ambiguity and insurance decisions. *American Economic Review Papers and Proceedings* 75, 386–390.
- Hogarth, R., Kunreuther, H., 1989. Risk ambiguity and insurance. *Journal of Risk and Uncertainty* 2, 5–35.
- Hogarth, R., Kunreuther, H., 1992. Pricing insurance and warranties: ambiguity and correlated risk. *Geneva Papers on Risk and Insurance Theory* 17, 35–60.
- Horie, M., 2007. A General Update Rule for Convex Capacities.** Working Papers 644, Kyoto University, Institute of Economic Research, November 2007.
- Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., Camerer, C., 2005. Neural systems responding to degrees of uncertainty in human decision-making. *Science* 310, 1680–1683.

- Huettel, S.A., Stowe, C., Gordon, E., Warner, B., Platt, M., 2006. Neural signatures of economic preferences for risk and ambiguity. *Neuron* 49, 765–775.
- Jaffray, J.-Y., 1992. Bayesian updating and belief functions. *IEEE Transactions on Systems, Man and Cybernetics* 22, 1144–1152.
- Jaffray, J.-Y., 1994. Dynamic decision making with belief functions. In: Yaeger, R.R., et al. (Eds.), *Advances in the Dempster-Shafer Theory of Evidence*. Wiley.
- Kahn, B., Sarin, R., 1988. Modeling ambiguity in decisions under uncertainty. *Journal of Consumer Research* 15, 265–272.
- Kelsey, D., Quiggin, J., 1992. Theories of choice under ignorance and uncertainty. *Journal of Economic Surveys* 6, 133–153.
- Keynes, J., 1921. *A Treatise on Probability*. Macmillan and Co, London.
- Kischka, P., Puppe, C., 1992. Decisions under risk and uncertainty: a survey of recent developments. *Mathematical Methods of Operations Research* 36, 125–147.
- Klibanoff, P., Marinacci, M., Mukerji, S., 2005. A smooth model of decision making under ambiguity. *Econometrica* 73, 1849–1892.
- Klibanoff, P., Marinacci, M., Mukerji, S., 2009. Recursive smooth ambiguity preferences. *Journal of Economic Theory* 144, 930–976.
- Klibanoff, P., Mukerji, S., Seo, K., 2011. *Relevance and Symmetry*. Discussion Paper N.539, Oxford University.
- Knight, F., 1921. *Risk, Uncertainty, and Profit*. Houghton, Mifflin, New York.
- Koch, C., Shunk, D., 2013. Limiting liability? – Risk and ambiguity attitudes under real losses. *SBR* 65, 54–75.
- Kopylov, I., 2006. *A Parametric Model of Ambiguity Hedging*. manuscript, University of California At Irvine.
- Kunreuther, H., 1989. The role of actuaries and underwriters in insuring ambiguous risks. *Risk Analysis* 9, 319–328.
- Kunreuther, H., Hogarth, R., 1992. How does ambiguity affect insurance decisions? In: Dionne, G. (Ed.), *Contributions To Insurance Economics*. Springer, New York.
- Kunreuther, H., Meszaros, J., Hogarth, R., Spranca, M., 1995. Ambiguity and underwriter decision processes. *Journal of Economic Behavior and Organization* 26, 337–352.
- Laplace, P., 1814. *Théorie analytique des probabilités*. Courcier, Paris.
- Larson, J., 1980. Exploring the external validity of a subjectively weighted utility model of decision making. *Organizational Behavior and Human Performance* 26, 293–304.
- Lehrer, E., 2007. *A Comment on an Example of Machina*. manuscript, Tel-Aviv University.
- Lehrer, E., 2009. A new integral for capacities. *Economic Theory* 39, 157–176.
- Lehrer, E., 2012. Partially-specified probabilities: decisions and games. *American Economic Journal: Microeconomics* 4, 70–100.
- LeRoy, S., Singell, L., 1987. Knight on risk and uncertainty. *Journal of Political Economy* 95, 394–406.
- L'Haridon, O., Placido, L., 2010. Betting on Machina's reflection example: an experiment on ambiguity. *Theory and Decision* 69, 375–393.
- Li, J., 2011. *Preferences for Information and Ambiguity*. manuscript, University of California, Berkeley.
- Maccheroni, F., Marinacci, M., Rustichini, A., 2006a. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74, 1447–1498.
- Maccheroni, F., Marinacci, M., Rustichini, A., 2006b. Dynamic variational preferences. *Journal of Economic Theory* 128, 4–44.
- MacCrimmon, K., 1965. *An Experimental Study of the Decision Making Behavior of Business Executives*. University of California, Los Angeles, Doctoral Dissertation.
- MacCrimmon, K., 1968. Descriptive and normative implications of the decision-theory postulates. In: Borch, K., Mossin, J. (Eds.), *Risk and Uncertainty: Proceedings of a Conference Held by the International Economic Association*. Macmillan and Co, London.
- MacCrimmon, K., Larsson, S., 1979. Utility theory: axioms versus 'paradoxes'. In: Allais, M., Hagen, G.M., (Eds.), *Expected Utility Hypotheses and the Allais Paradox*. D. Reidel, Dordrecht.
- Machina, M., 1987. Choice under uncertainty: problems solved and unsolved. *Journal of Economic Perspectives* 1, 121–154.

- Machina, M., 1989. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature* 27, 1622–1668.
- Machina, M., 2004. Almost-objective uncertainty. *Economic Theory* 24, 1–54.
- Machina, M., 2009. Risk, ambiguity, and the rank-dependence axioms. *American Economic Review* 99, 385–392.
- Machina, M., 2013. *Ambiguity Aversion with Three or More Outcomes*. manuscript, University of California, San Diego.
- Machina, M., Schmeidler, D., 1992. A more robust definition of subjective probability. *Econometrica* 60, 745–780.
- Machina, M., Schmeidler, D., 1995. Bayes without Bernoulli: simple conditions for probabilistically sophisticated choice. *Journal of Economic Theory* 67, 106–128.
- Maffioletti, A., Santori, M., 2005. Do trade union leaders violate subjective expected utility? Some insights from experimental data. *Theory and Decision* 59, 207–253.
- Malinvaud, E., 1952. Note on von Neumann-Morgenstern's strong independence axiom. *Econometrica* 20, 679.
- Marschak, J., 1950. Rational behavior, uncertain prospects, and measurable utility. *Econometrica* 18, 111–141 (Errata, *Econometrica* 18, 312).
- Milnor, J., 1954. Games against nature. In: Thrall, R., Coombs, C., Davis, R. (Eds.), *Decision Processes*. John Wiley & Sons, New York.
- Mukerji, S., 2000. A survey of some applications of the idea of ambiguity aversion in economics. *International Journal of Approximate Reasoning* 24, 221–234.
- Nakamura, N., 1990. Subjective expected utility with non-additive probabilities on finite state spaces. *Journal of Economic Theory* 51, 346–366.
- Nau, R., 2006. Uncertainty aversion with second-order utilities and probabilities. *Management Science* 52, 136–145.
- Nehring, K., 2001. *Ambiguity in the Context of Probabilistic Beliefs*. manuscript, University of California, Davis.
- Neyman, J., Pearson, E., 1933. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society, Series A* 231, 289–337.
- Olszewski, W., 2007. Preferences over sets of lotteries. *Review of Economic Studies* 74, 567–595.
- Ortoleva, P., 2010. Status quo bias, multiple priors and uncertainty aversion. *Games and Economic Behavior* 69, 411–424.
- Ozdenoren, E., Peck, J., 2008. Ambiguity aversion, games against nature, and dynamic consistency. *Games and Economic Behavior* 62, 106–115.
- Pires, C., 2002. A rule for updating ambiguous beliefs. *Theory and Decision* 53, 137–152.
- Poincaré, H., 1912. *Calcul des Probabilités*, second edn. Gauthiers-Villars, Paris.
- Pratt, J., 1964. Risk aversion in the small and in the large. *Econometrica* 32, 122–136.
- Pratt, J., Raiffa, H., Schlaifer, R., 1964. The foundations of decisions under uncertainty: an elementary exposition. *Journal of the American Statistical Association* 59, 353–375.
- Quiggin, J., 1982. A theory of anticipated utility. *Journal of Economic Behavior and Organization* 3, 323–343.
- Raiffa, H., 1961. Risk, ambiguity, and the savage axioms: comment. *Quarterly Journal of Economics* 75, 690–694.
- Ramsey, F., 1926. Truth and probability. In: Ramsey, F. (Ed.), *Foundations of Mathematics and other Logical Essays*. K. Paul, Trench, Trubner and Co, London (1931).
- Ritov, I., Baron, J., 1990. Reluctance to vaccinate: omission bias and ambiguity. *Journal of Behavioral Decision Making* 3, 263–277.
- Roberts, H., 1963. Risk, ambiguity, and the savage axioms: comment. *Quarterly Journal of Economics* 77, 327–336.
- Rothschild, M., Stiglitz, J., 1970. Increasing risk: I. a definition. *Journal of Economic Theory* 2, 225–243.
- Samuelson, P., 1952. Probability, utility, and the independence axiom. *Econometrica* 20, 670–678.
- Sarin, R., Wakker, P., 1992. A Simple axiomatization of nonadditive expected utility. *Econometrica*, 1255–1272.

- Sarin, R., Winkler, R., 1992. Ambiguity and decision modeling: a preference-based approach. *Journal of Risk and Uncertainty* 5, 389–407.
- Savage, L.J., 1950. The role of personal probability in statistics. *Econometrica* 18, 183–184.
- Savage, L., 1954. *The Foundations of Statistics*. John Wiley & Sons, New York. Revised and Enlarged Edition, New York: Dover Publications, 1972..
- Savage, L., 1963. Leonard J. Savage Papers. Yale University Library Manuscripts and Archives, Collection Number 695.
- Schmeidler, D., 1986. Integral representation without additivity. *Proceedings of the American Mathematical Society* 97, 255–261.
- Schmeidler, D., 1989. Subjective probability and expected utility without additivity. *Econometrica* 57, 571–587.
- Segal, U., 1987a. The ellberg paradox and risk aversion: an anticipated utility approach. *International Economic Review* 28, 175–202.
- Segal, U., 1987b. Some remarks on quiggin's anticipated utility. *Journal of Economic Behavior and Organization* 8, 145–154.
- Segal, U., 1990. Two-stage lotteries without the reduction axiom. *Econometrica* 58, 349–377.
- Seo, K., 2009. Ambiguity and second-order belief. *Econometrica* 77, 1575–1605.
- Shackle, G., 1949a. *Expectation in Economics*. Cambridge University Press, Cambridge.
- Shackle, G., 1949b. A non-additive measure of uncertainty. *Review of Economic Studies* 17, 70–74.
- Shafer, G., 1976. *A Mathematical Theory of Evidence*, vol. 1. Princeton University Press, Princeton.
- Shafer, G., 1978. Non-additive probabilities in the work of Bernoulli and Lambert. *Archive for History of the Exact Sciences* 19, 309–370.
- Siniscalchi, M., 2006. A behavioral characterization of plausible priors. *Journal of Economic Theory* 128, 91–135.
- Siniscalchi, M., 2008a. Ambiguity and ambiguity aversion. In: Durlauf, S., Blume, L. (Eds.), *The New Palgrave Dictionary of Economics*, second edn. Macmillan, London.
- Siniscalchi, M., 2008b. *Machina's Reflection Example and VEU Preferences: A Very Short Note*. manuscript, Northwestern University.
- Siniscalchi, M., 2009a. Vector expected utility and attitudes toward variation. *Econometrica* 77, 801–855.
- Siniscalchi, M., 2009b. Two out of three ain't bad: a comment on 'the ambiguity aversion literature: a critical assessment'. *Economics and Philosophy* 25, 335–356.
- Siniscalchi, M., 2010. *Recursive Vector Expected Utility*, manuscript. Northwestern University.
- Siniscalchi, M., 2011. Dynamic choice under ambiguity. *Theoretical Economics* 6, 379–421.
- Sinn, H.-W., 1980. A rehabilitation of the principle of sufficient reason. *Quarterly Journal of Economics* 94, 493–506.
- Skiadas, C., 2013. Scale-invariant uncertainty-averse preferences and source-dependent constant relative risk aversion. *Theoretical Economics* 8, 59–93.
- Slovic, P., Tversky, A., 1974. Who accepts savage's axiom? *Behavioral Science* 19, 368–373.
- Smith, V., 1969. Measuring nonmonetary utilities in uncertain choices: the Ellsberg urn. *Quarterly Journal of Economics* 83, 324–329.
- Starmer, C., 2000. Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *Journal of Economic Literature* 38, 332–382.
- Stigler, S., 1986. *The History of Statistics: The Measurement of Uncertainty Before 1900*. Harvard University Press, Cambridge, MA.
- Strotz, R.H., 1956. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23, 165–180.
- Strzalecki, T., 2011. Axiomatic foundations of multiplier preferences. *Econometrica* 79, 47–73.
- Sugden, R., 1986. New developments in the theory of choice under uncertainty. *Bulletin of Economic Research* 38, 1–24.
- Turvey, R., 1949. Note on 'expectation in economics'. *Economica* 16, 336–338.
- Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: heuristics and biases. *Science* 185, 1124–1131.
- Venn, J., 1866. *The Logic of Chance*. Macmillan and Co, London. 2nd Ed. 1876. 3rd Ed. 1888.

- Viscusi, W., 1989. Prospective reference theory: toward an explanation of the paradoxes. *Journal of Risk and Uncertainty* 2, 235–264.
- Viscusi, W.K., Chesson, H., 1999. Hopes and fears: the conflicting effects of risk ambiguity. *Journal of Risk and Uncertainty* 47, 153–178.
- Viscusi, W., Magat, W., 1992. Bayesian decisions with ambiguous belief aversion. *Journal of Risk and Uncertainty* 5, 371–387.
- Viscusi, W., Magat, W., Huber, J., 1991. Communication of ambiguous risk information. *Theory and Decision* 31, 159–173.
- von Kries, J., 1886. *Die Principien der Wahrscheinlichkeits-Rechnung: Eine logische Untersuchung*. Mohr, Tübingen.
- von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton (2nd Ed. 1947; 3rd Ed. 1953).
- Wakker, P., 1989. *Additive Representations of Preferences: A New Foundation of Decision Analysis*. Kluwer Academic Publishers, Dordrecht, Holland.
- Wakker, P., 2010. *Prospect Theory: For Risk and Ambiguity*. Cambridge University Press, Cambridge, UK.
- Wakker, P., 2013. *Annotated Bibliography*. Available from: <<http://people.few.eur.nl/wakker/refs/webrfnrcs.doc>>.
- Wald, A., 1939. Contributions to the theory of statistical estimation and testing hypotheses. *Annals of Mathematical Statistics* 10, 299–326.
- Wald, A., 1950. Some recent results in the theory of statistical decision functions. *Econometrica* 18, 182–183.
- Walley, P., 1991. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London.
- Wang, T., 2003. *A Class of Multi-Prior Preferences*. manuscript, Sauder School of Business, University of British Columbia.
- Weber, E.U., Johnson, E., 2008. Decisions under uncertainty: psychological, economic, and neuroeconomic explanations of risk preference. In: Glimcher, P., Camerer, C., Fehr, E., Poldrack, R. (Eds.), *Neuroeconomics: Decision Making and the Brain*. Elsevier, New York.
- Willham, C., Christensen-Szalanski, J., 1993. Ambiguity and liability negotiations: the effects of the negotiators' role and the sensitivity zone. *Organizational Behavior and Human Decision Processes* 54, 277–298.
- Yaari, M.E., 1969. Some remarks on measures of risk aversion and on their uses. *Journal of Economic Theory* 1, 315–329.
- Yaari, M.E., 1987. The dual theory of choice under risk. *Econometrica* 55, 95–115.
- Yang, C.-L., Yao, L., 2011. *Ellsberg Paradox and Second-order Preference Theories on Ambiguity: Some New Experimental Evidence*. manuscript, Academia Sinica.
- Yates, F.J., Zukowski, L.G., 1976. Characterization of ambiguity in decision making. *Behavioral Science* 21, 19–25.
- Zhang, J., 2002. Subjective ambiguity, expected utility and choquet expected utility. *Economic Theory* 20, 159–181.