

Expected Utility Theory

Simon Grant
Rice University

Timothy Van Zandt
INSEAD

22 November 2007

This is a chapter for the forthcoming *Handbook of Rational and Social Choice*, Paul Anand, Prasanta Pattanaik, and Clemens Puppe, eds., Oxford University Press, 2008. We review classic normative expected utility theory. Our goal is to frame the subsequent chapters (which consider more modern extensions to and deviations from this classic theory) in a way that is accessible to the non-specialist but also useful to the specialist. We start from scratch with a revealed preference approach to the existence of a utility function. We then present the mathematical structure of additive and linear utility representations and their axiomatizations, in the context of abstract choice theory and using intertemporal choice as a source of examples. We are thus able to focus on this mathematical structure without the interference of the specific interpretation and notation for decision under uncertainty. Furthermore, this approach allows us to focus on the interpretation of the axioms when we turn to decision under uncertainty.

Contents

1. Introduction	1
2. Descriptive, prescriptive, and normative theories	2
3. A review of choice, preferences, and utility	3
4. Hidden assumption: Consequentialism	3
5. Weak axiom of revealed preference	4
6. From preferences to choice rules	6
7. Infinite choice sets and continuity assumptions	7
8. Additive utility	8
9. Linear preferences	12
10. Cardinal uniformity across factors	15
11. States of nature and acts	20
12. Sure-thing principle and additivity	22
13. Lotteries	27
14. Subjective expected utility without objective probabilities	31
15. Subjective expected utility with objective probabilities	33
16. Conclusion	36
References	40

Acknowledgments: We are deeply indebted to Peter Wakker for providing detailed comments on two earlier drafts and helping shape this chapter. We also benefited from the feedback of Paul Anand and Clemens Puppe, from discussions with Enrico Diecidue, and from corrections by Philippe Delqu   and Alminas Zaldokas.

Email addresses: sgrant@rice.edu (Grant); tvz@insead.edu (Van Zandt).

1. Introduction

This handbook is a modern look at decision theory and social choice. It emphasizes recent developments that go beyond the classic theory of utility and expected utility. However, a sensible starting point is the classical theory, a benchmark that will frame departures considered in subsequent chapters. Furthermore, it remains the prominent workhorse for applied economics.

This chapter presents the main features of that classical theory. The topic is broad and has been treated so much elsewhere that we do not attempt to be technically comprehensive. Instead, we opt for a succinct and pedagogical treatment that will be accessible to nonspecialists yet still useful to the specialist for framing the subsequent chapters.

The theory reviewed here has had a long and complicated development. Expected utility dates back at least to Bernoulli (1738). As a resolution to the famous St. Petersburg paradox, he used a logarithmic utility index defined over wealth to compute a finite price for a gamble with an unbounded expected value. Even the modern foundations are a half-century old, with the development of an axiomatic choice-theoretic foundation for expected utility occurring roughly from the mid-1920s to the early 1960s.

We neither structure the chapter around such historical development nor attempt to explain it; a proper treatment would be too lengthy and would distract from the content of this chapter, whereas a succinct treatment would have too many unfair omissions. Instead, our goal is to clarify the interpretation and mathematical role of the various axioms that form the theory. We have the benefit of hindsight and of our freedom to deviate from the chronology of the theory's development.

Rather than delve immediately into choice under uncertainty, we start from scratch with basic preference theory (Sections 2–7). Not all of this volume relates to decision under uncertainty, but our other motive for reviewing this theory is to emphasize the assumptions that are already implicit when one represents choices over uncertain prospects by a complete and transitive preference relation.

The fundamental mathematical properties of expected utility representations are additivity and linearity. Our main expository innovation is to remain in the general choice setting a little longer (Sections 8–10) in order to explore the independence assumptions that lead to such structure *without reference to decision under uncertainty*. We can thereby (a) clarify the mathematics without the baggage of the uncertainty framework and (b) discuss the axioms in other contexts, such as multi-attribute decision problems and intertemporal choice, so that their significance in decision under uncertainty is then better understood.

We thus do not reach decision under uncertainty until the middle of the chapter. We first introduce the “states of the world” (Savage) framework for representing uncertainty in Section 11. This precedes our treatment of the lotteries representation of objective uncertainty (Section 13), in a reversal of the usual (and historical) order, because the states framework is useful and intuitive whether uncertainty is subjective or objective and because we refer to it when interpreting the axioms of the lotteries model. Overall, we proceed as follows. After presenting the states framework in Section 11, we take our characterization of preferences

and utility as far as an independence axiom and state-dependent expected utility in Section 12. We then turn to expected utility for lotteries (objective uncertainty) in Section 13. In Sections 14 and 15, we return to the full subjective expected utility representation in the states model with state-independent utility.

As noted, we emphasize the link between the independence axioms in consumer theory, in expected utility for objective lotteries, and in expected utility under subjective uncertainty. Fishburn and Wakker (1995) provide a comprehensive study of how these different versions arose and evolved historically. The way in which we motivate the axioms—as following from dynamic consistency and consequentialism—is analogous to the treatment in Hammond (1988).

Our decision to study additive and linear utility before any applications to decision under uncertainty has the advantages we claimed above. It comes at two costs. First, those familiar with expected utility theory may be disconcerted to see, in the first half of the chapter, familiar axioms and theorems with only passing references to the seminal work in expected utility theory from which they originated. However, such omissions are rectified in the second half of the chapter and, with this warning, the reader should suffer no ill consequences. The second cost is that we have to invent “context-free” names for several old axioms from expected utility theory.

2. Descriptive, prescriptive, and normative theories

Decision theory has two goals: To *describe* how agents *do* make decisions (descriptive decision theory) and to *prescribe* how agents *should* make decisions (prescriptive decision theory). As in any theoretical modeling, decision theory balances accuracy and simplicity. A prescriptive decision theory that is too complicated to learn or implement is hardly useful for a decision maker. A descriptive theory should be simple because it is meant to be a framework that organizes and unifies a variety of situations and behavior, because it should be tractable enough to derive conclusions, and because we may need to estimate the parameters of the theory from a limited amount of data.

A third branch of decision theory, normative decision theory, tries to describe how a hypothetical, infinitely intelligent being would make decisions. This may sound speculative and impractical but it provides important foundations for descriptive and prescriptive theories. A normative theory is inherently simpler than an overtly descriptive or prescriptive theory because it need not concern itself with such complications as (a) errors or forgetting and (b) the heterogeneity of the intelligence and experience of decision makers. There are only a few ways to be perfect, but many ways to be imperfect!

Simplicity is good, but how does a normative theory serve the goals of descriptive or prescriptive theory? Humans are goal oriented and work hard to pursue these goals. Any descriptive theory should capture this first-order consideration, and a normative model is a powerful and parsimonious way to do so. To develop a prescriptive theory that helps mortal humans (with all their limitations) make decisions, it is useful to understand how unboundedly rational beings would make decisions.

This chapter develops classical rational choice theory, particularly expected utility theory, as a normative model. We leave it to subsequent chapters to evaluate how well it serves as a descriptive or prescriptive theory.

3. A review of choice, preferences, and utility

Before tackling decision making under uncertainty, we study the theory of choice without uncertainty. Let's call our decision maker "Anna".

In most of this chapter, we limit attention to *static choice*. This means we consider a single decision that Anna makes. We thus suppress the fact that Anna may anticipate having to make choices in the future and that current decisions are intertwined with such future decisions. However, we do not thereby suppress time: the objects of choice could be time paths of consumption.

We ignore much of the information processing that Anna must do in order to make a decision. The only input to Anna's problem is the set of possible alternatives, which we allow to vary. There is a fixed set X containing all potential alternatives, and Anna is presented with a set $A \subset X$ from which she must choose, called a *choice set*. For example, X is the set of all consumption bundles and A is the budget set, which depends on prices and the agent's wealth. Alternatively, X could be the set of potential presidential candidates and A the set of candidates on the ballot. We want a model of what Anna would choose from each set of feasible alternatives.

Assume that X is finite and that every nonempty subset A of X is a potential feasible set. For each set $A \subset X$ of feasible alternatives, let $C(A)$ be the elements of A that Anna might choose from A . She must always choose something, which means that $C(A)$ is nonempty, but $C(A)$ may contain more than one item because of indifference. Call $C(\cdot)$ her *choice rule*. Our game is to find some conditions that lead to a simple representation of choice. The representation should allow one to derive qualitative conclusions in models without knowing specific choices and should have few parameters, all of which could be estimated empirically.

There are two complementary approaches. One is to start with a choice rule as a primitive. From this "empirical" object, one derives (revealed) preferences. The other approach is to start with preferences as the primitive. From these, one derives the choice rule. We follow the first (revealed preference) approach.

4. Hidden assumption: Consequentialism

We will be as explicit as space permits about hidden assumptions. For example, we have already outlined the static nature of the model. Another hidden assumption is consequentialism: Anna cares about *consequences*—not how consequences are achieved.

In the decision model we have developed so far, "consequences" refers to the alternatives and "consequentialism" means that Anna does not care how she ends up facing a particular choice set $A \subset X$ or how she chooses from A . For example, in consumer theory, one might suppose that Anna cares only about consumption

and not about the prices that determine which consumption bundles are affordable. In decision making under uncertainty, we may suppose that Anna cares only about probabilities of different outcomes and not about how these probabilities are generated (e.g., about whether uncertainty is resolved in one step or instead in a multi-stage lottery).

We can redefine consequences in order to circumvent any particular violation of consequentialism. If Anna cares about prices, then we define a consequence not only by how much she consumes of different goods but also by the prices she faces (we have thus redefined the set X). If Anna cares about the stages at which uncertainty unfolds, then a consequence can be defined to include such information.

However, if we proceed this way without constraint, then—in our abstract model—a consequence would be defined by not only the element of X that Anna ends up with but also by the set A from which she was able to select it. Hence, there would be no link or consistency conditions between the choices Anna might make from $A \subset X$ and those she might make from $A' \subset X$. We would have a vacuous theory. Instead, we must be able to define consequences in a sufficiently restrictive way that we can envision Anna being presented with different sets of consequences and caring only about those consequences, not about how those sets were generated.

From the empirical side of the theory, there is another hidden assumption. We are supposedly studying a single decision that Anna makes. However, for that decision she will end up facing a single choice set $A \subset X$. We will never be able to observe an inconsistency between $C(A)$ and $C(A')$ unless one of the following two conditions is satisfied.

1. When presented with different *hypothetical* choice sets, Anna can report the choices she would make.
2. Anna faces different instances of this problem at different times, and how she chooses from any choice set does not vary from one instance to the other.

Thus, we assume that one of these two conditions holds.

5. Weak axiom of revealed preference

Let x and y belong to X . Then x is *revealed weakly preferred* to y if $x \in C(A)$ for some $A \subset X$ containing x and y (i.e., if y is available but x may be chosen). If also $y \notin C(A)$ then we say that x is *revealed preferred* to y or, if we want to emphasize that the preference is not weak, that x is *revealed strictly preferred* to y .

The *weak axiom of revealed preference* (WARP) states that if x is revealed weakly preferred to y , then y is not revealed preferred to x . This is defined explicitly in Axiom 1.

AXIOM 1 (WARP). Let x and y belong to X , and let A and B be subsets of X containing x and y . If $x \in C(A)$ and $y \in C(B)$, then $x \in C(B)$.

WARP is a natural axiom of a normative theory, but it is at best an approx-

imation for descriptive or prescriptive theories. Especially when choice sets are large and complex, achieving such consistency is difficult.

This one consistency condition gets us far. First, it implies that the choice rule can be summarized by or deduced from more limited information: the binary choices or *preferences*. To make this statement more precise, we first define, for $x, y \in X$,

$$\begin{aligned} x \succcurlyeq y & \text{ if } x \in C(\{x, y\}) && (x \text{ is } \textit{weakly preferred} \text{ to } y); \\ x \succ y & \text{ if } x \succcurlyeq y \text{ but not } y \succcurlyeq x && (x \text{ is } \textit{strictly preferred} \text{ to } y); \\ x \sim y & \text{ if } x \succcurlyeq y \text{ and } y \succcurlyeq x && (x \text{ is } \textit{indifferent} \text{ to } y). \end{aligned}$$

The binary relations \succcurlyeq , \succ , and \sim are called the (weak) preference, strict preference, and indifference relations, respectively.

DEFINITION 1. Let \succcurlyeq be the preference relation defined for a choice rule $C(\cdot)$ as in the previous paragraph. Then the choice rule satisfies *preference maximization* if, for every $A \subset X$ and $x \in A$,

$$x \in C(A) \iff x \succcurlyeq y \quad \forall y \in A.$$

In words, choices from large sets are consistent with binary choices. If we know Anna's binary choices (preferences), then we can derive $C(\cdot)$. Preference maximization implies a considerable savings in the amount of information required in order to know $C(\cdot)$.

It is also useful for Anna's preferences to satisfy some consistency conditions themselves.

DEFINITION 2. A binary relation \succcurlyeq is a *weak order* if it satisfies the following conditions.

1. (*Completeness*) For all $x, y \in X$, we have $x \succcurlyeq y$ or $y \succcurlyeq x$ (or both).
2. (*Transitivity*) For all $x, y, z \in X$, if $x \succcurlyeq y$ and $y \succcurlyeq z$ then $x \succcurlyeq z$.

Completeness of the preference relation follows from the assumption that C has nonempty values. Transitivity is the important consistency condition; it follows from WARP.

PROPOSITION 1. *The choice rule $C(\cdot)$ satisfies WARP if and only if (a) it satisfies preference maximization and (b) the preference relation is complete and transitive.*

Proof. Samuelson (1938) defined WARP in the context of consumer theory (budget sets) and single-valued demand. Arrow (1959) defined WARP (as we did) for general choice rules—that is, for abstract choice sets and possibly multi-valued choices. Proposition 1 is implied by Theorems 2 and 3 in Arrow (1959). \square

One advantage of complete and transitive preferences is that they can be represented by a utility function.

DEFINITION 3. Let \succsim be a preference relation on X and let $U: X \rightarrow \mathbb{R}$ be a function. Then U is a *utility representation* of \succsim , and \succsim is *represented by the utility function* U , if for all $x, y, \in X$ we have

$$x \succsim y \iff U(x) \geq U(y).$$

PROPOSITION 2. A preference relation \succsim on X is complete and transitive if and only if it has a utility representation.

Proof. The utility representation is easily constructed recursively. See Birkhoff (1948, Thm. 1, p. 31) for such a proof for countable X . \square

Thus, WARP also means that Anna makes choices *as if* she maximized a utility function, as stated in the next corollary.

COROLLARY 1. The choice rule $C(\cdot)$ satisfies WARP if and only if there is a utility function $U: X \rightarrow \mathbb{R}$ such that

$$C(A) = \arg \max_{x \in A} U(x)$$

for all $A \subset X$.

If $U: X \rightarrow \mathbb{R}$ is a utility representation of a preference relation \succsim on X and if $f: \mathbb{R} \rightarrow \mathbb{R}$ is any strictly increasing function, then $f \circ U: X \rightarrow \mathbb{R}$ is also a representation of \succsim . (The composition $f \circ U$ is called a *monotone transformation* of U). Therefore, the magnitudes of the utility values of a representation have no particular meaning—they only define the ordinal relationship \succsim .

6. From preferences to choice rules

Suppose instead that we let Anna's preference relation \succsim be the primitive and then derive choices from this relation. (The primitive is thus any binary relation on X , which attains meaning as a preference relation through the subsequent use made of it in the model.) In this approach, preference maximization becomes an *axiom* that defines the choice rule: For $A \subset X$,

$$C(A) = \{x \in A \mid x \succsim y \forall y \in A\}.$$

Given this axiom, the standard way to proceed is by assuming that preferences are complete and transitive and then concluding that (a) the choice rule $C(\cdot)$ satisfies WARP and (b) preferences have a utility representation.

We highlight the axiom of preference maximization—that choices from large sets be consistent with binary choices—because it is a critical consistency condition. If one observed a violation of WARP or a violation of transitivity, then it would be reasonable to doubt preference maximization. Retaining it and merely tweaking the axioms on preferences could be a poor way to develop a theory that encompasses such empirical violations.

7. Infinite choice sets and continuity assumptions

Infinite choice sets require two modifications to the theories described so far. One such modification is that the choice rule might not be defined on all subsets of X : it has a restricted domain. For example, if $X = \mathbb{R}_+$ represents monetary prizes and if Anna would always choose the higher of two monetary prizes, then preference maximization would imply that Anna's choice from all of X or from $[0, 1)$ is not well-defined.

This should be viewed as a nonsubstantive technical restriction. In practice, it is fair to say that Anna would always make a decision. (Not announcing a choice also represents a decision, since there must be some default outcome that occurs.) The kind of choice sets that are suppressed from the domain of C typically either are unrealistic ("you can have as much money as you want"), arise only because an inherently discrete problem has been modeled as a continuous approximation ("you can have any amount of money in $[0, 1)$ "), or would lead to well-defined choices if certain aspects of the problem were fully modeled (e.g., if the time it took to get finer and finer divisions of a good were taken into account).

The second modification of the theory, which may be needed only if the choice set is uncountable, is that there may be an important role for some kind of continuity assumption on the preference relation.

DEFINITION 4. Let X be endowed with a topology. A preference relation \succsim on X is *continuous* if $\{y \in X \mid y \succsim x\}$ and $\{y \in X \mid x \succsim y\}$ are closed for all $x \in X$.

One representation theorem on infinite choice sets is the following.

THEOREM 1. *Suppose that X is a separable metric space and that \succsim is continuous, complete, and transitive. Then \succsim has a continuous utility representation.*

Proof. See Theorem II in Debreu (1954). The actual topological assumption is more general: that the topology on X have a countable base. For a metric space, this assumption is equivalent to separability. See Debreu (1964) for a discussion and alternate proofs of variants on this result. \square

In Theorem 1, the continuity assumption is needed in order to obtain any utility representation at all, but it has the added benefit that the theorem shows the existence of a continuous representation. In the approach that starts at preferences and derives choices, continuity assumptions are used to ensure that choices are well-defined at least on compact sets. Continuity of the representation may be important in further applications.

However, it would be surprising if continuity of preferences were required for the mere existence of a (not necessarily continuous) utility representation. In fact, there are nontopological variants of Theorem 1 (and of many other theorems in this chapter) that use an "order-theoretic" or "algebraic" approach. We present one example.

Define a subset Z of (X, \succsim) (or of any weakly-ordered space) to be *order-dense* in (X, \succsim) if, for all $x, y \in X$ such that $x \succ y$, there exists $z \in Z$ such that

$x \succ z \succ y$. For any $Y \subset \mathbb{R}$, the linearly-ordered space (Y, \geq) has a *countable* order-dense subset. One can think of a utility representation as embedding the ordered space (X, \succ) into (\mathbb{R}, \geq) . Therefore, (X, \succ) must also have a countable order-dense subset. This necessary condition is also sufficient for the existence of a utility representation.

THEOREM 2. *If \succ is complete and transitive and contains a countable order-dense subset, then \succ has a utility representation.*

Proof. This is Lemma II in Debreu (1954). It is closely related to results for linear orders by Cantor (1895) and Birkhoff (1948, Thm. 2, p. 32). (See a complete proof of Birkhoff’s theorem in Krantz et al. (1971, Thm. 2.2).) One proves Theorem 2 by studying the induced linear order over equivalence classes of (X, \succ) and thereby reducing it to the case of linear orders. This is made explicit in Theorem 3.1 of Fishburn (1970).

We provide a warning to the reader. The term “order dense” has several similar definitions in the mathematics and mathematical economics literature. For example, for a linearly-ordered space (Y, \geq) , Birkhoff (1948) defines $Z \subset Y$ to be *order dense* in (Y, \geq) if, for all $x, y \in Y \setminus Z$ such that $x > y$, there is $z \in Z$ such that $x > z > y$. This definition is equivalent to ours for a linearly-ordered set but not for a weakly-ordered set. \square

Theorems 1 and 2 are linked by the following fact: If X is a separable metric space and \succ is continuous, then (X, \succ) has a countable order-dense subset.

An example of preferences that have no countable order-dense subset and hence have no utility representation are lexicographic preferences on $\mathbb{R}_+^2: (x_1, x_2) \succ (y_1, y_2)$ if and only if either $x_1 > y_1$ or $x_1 = y_1$ and $y_1 > y_2$. Any order-dense subset of (\mathbb{R}_+^2, \succ) is not countable because it must contain an element for every value of the first coordinate.

Apart from this brief aside, we stick to the topological approach throughout this chapter. The topological assumptions tend to be both easier to state and more familiar to the reader than are the algebraic assumptions. The interested reader may consult Krantz et al. (1971) and Fishburn (1970) for further development of the algebraic approach.

8. Additive utility

Maintained assumptions. The rest of this chapter is about the *additional* structure that preferences and their utility representation may have. Therefore, as we proceed and even as the set of alternatives varies, we avoid repetition by *maintaining the assumptions that choices satisfy preference maximization, that the symbol \succ denotes the preferences relation, and that \succ is complete and transitive.*

8.1. Overview

We study two types of structure that one might impose on preferences or utility: additivity and linearity. Why do we tackle these technical topics in the setting of abstract choice theory when this chapter is supposed to focus on choice under uncertainty? In decision under uncertainty, additivity is the most important structure imposed on preferences over state-contingent outcomes, and linearity is the most important assumption imposed on preferences over lotteries. We would like to understand the mathematics of these structures in the absence of a particular interpretation so that (a) the mathematical structure is clearer, (b) when we reach decision under uncertainty, we can focus on their interpretation rather than the math, and (c) we can better understand the relationship between expected utility theory and, say, intertemporal choice.

8.2. Additive separability

In this section, we assume that X is a product set: $X = X_1 \times \cdots \times X_J$. Each component $j = 1, \dots, J$ is called a *factor*. Here are some possible cases.

1. There are J goods; $(x_1, \dots, x_J) \in X$ is a consumption bundle; and $X_j = \mathbb{R}_+$ is the set of possible quantities of any one of the goods.
2. There is one good with multiple varieties, which are characterized by J attributes; a particular variety is given by a list $(x_1, \dots, x_J) \in X$ of attributes; and X_j is the set of possible values of attribute j .
3. There are J periods; $(x_1, \dots, x_J) \in X$ is a time path of consumption; and X_j is the set of possible consumption bundles in period j .
4. There are J states of nature; $(x_1, \dots, x_J) \in X$ is a state-contingent outcome; and X_j is the set of possible outcomes in state j .
5. There are J people; $(x_1, \dots, x_J) \in X$ is an allocation; and X_j is the set of possible consumption bundles of agent j . (The decision maker is an outside observer.)

Given preferences \succsim on X , we are interested in whether there is a representation of the form

$$U(x_1, \dots, x_J) = \sum_{j=1}^J u_j(x_j)$$

for some functions $u_j: X_j \rightarrow \mathbb{R}$. Such a representation is called *additive* or *additively separable*.

Suppose such a representation exists. Recall that any monotone transformation of U represents the same preferences. However, it need not be additive. The family of additive representations can be fairly narrow. Under a variety of assumptions, including those of Theorem 3 to follow, an additive representation is unique up to an affine transformation. That is, V is another additive representation if and only if $V = a + bU$, where $a \in \mathbb{R}$ and $b \in \mathbb{R}_{++}$. The magnitudes of the utilities of the different factors have some meaning because they are aggregated across factors to determine the overall ranking of alternatives.

For example, suppose $x^1, x^2, x^3, x^4 \in X$, with $x^1 \succ x^2$ and $x^3 \succ x^4$. Compare the “extra kick” from getting x^1 rather than x^2 with the extra kick of getting x^3

rather than x^4 and measure this comparison by the ratio

$$\frac{U(x^1) - U(x^2)}{U(x^3) - U(x^4)}. \quad (1)$$

This ratio is the same for any additive representation U , yet can have any positive value if we allow for nonadditive representations.

For this reason, additive utilities are often referred to as *cardinal* utilities. However, only ratios such as (1) are uniform across additive representations. Thus, additive utilities are not interpersonal cardinal scales as used in utilitarianism.

8.3. Joint independence

An immediate consequence of an additive representation is that preferences satisfy *joint independence*, meaning that “how one ranks what happens with some factors does not depend on what happens with the other factors”.

We introduce some notation to formalize joint independence. For a given set $K \subset \{1, \dots, J\}$ of factors, let $X_K \equiv \prod_{j \in K} X_j$. For a partition $\{K, L\}$ of $\{1, \dots, J\}$, we may write an element of X as (a, c) , where $a \in X_K$ denotes the values for the factors in K and $c \in X_L$ denotes the values for the factors in L . For the special case in which K is a singleton $\{j\}$, we may write (x_j, x_{-j}) , where $x_j \in X_j$ and $x_{-j} \in X_{-j} \equiv \prod_{k \neq j} X_k$.

AXIOM 2 (Joint independence). For all partitions $\{K, L\}$ of $\{1, \dots, J\}$, for all $a, b \in X_K$, and for all $c, d \in X_L$,

$$(a, c) \succ (b, c) \iff (a, d) \succ (b, d).$$

Joint independence (a term used by Krantz et al. 1971, p. 339) has also been called *conjoint independence*, *strong separability* (Strotz 1959), *independence among factors* (Fishburn 1970, Sec. 4.1; Debreu 1960), and *coordinate independence* (Wakker 1989). It was called the *sure-thing principle* by Savage (1954) for a model in which the factors are states.

We mentioned intertemporal choice as a setting in which one may want to work with additive representations. Note how strong the joint independence condition can be in that setting. It implies that how Anna prefers to allocate consumption between periods 2 and 3 does not depend on how much she consumed in period 1. However, after a shopping and eating binge on Tuesday, Anna might well prefer to take it easy on Wednesday and defer some consumption until Thursday. Joint independence also implies that Anna’s ranking of today’s meal options does not depend on what she ate yesterday and that whether she wants to talk to her boyfriend today does not depend on whether they had a fight yesterday.

Yet additive utility is a mainstay of intertemporal models in economics. The assumption becomes a better approximation when outcomes are measured at a more aggregate level, as is typical in such models. Suppose, then, that the model does not include such details as meal options and social relationships, addressing instead one-dimensional aggregate consumption. Then joint independence

in a two-period model is satisfied as long as Anna's preferences are monotone (though we will see that, in this case, the assumption is not enough to guarantee an additive utility representation). With more than two time periods, joint independence implies that Anna's desired allocation between periods 2 and 3 does not depend on period-1 consumption; this assumption is more plausible when a period is one year rather than one day.

8.4. A representation theorem

Joint independence is the main substantive assumption needed to obtain an additive representation. Here is one such representation theorem.

THEOREM 3. *Assume $J \geq 3$ and that the following hold.*

1. X_j is connected for all $j = 1, \dots, J$.
2. \succsim is continuous.
3. \succsim satisfies joint independence.
4. Each factor is essential: for all j , there exist $x_j, x'_j \in X_j$ and $x_{-j} \in X_{-j}$ such that $(x_j, x_{-j}) \succ (x'_j, x_{-j})$.

Then \succsim has an additive representation that is continuous.

Proof. This result is similar to Debreu (1960, Thm. 3), except that the latter also assumes that each X_j is separable. Krantz et al. (1971, Thm. 13, Sec. 6.11) provides an algebraic approach to such a representation; the topological assumptions are replaced by algebraic conditions on the preferences and the continuity of the additive representation is not derived. Their Theorem 14 then shows that our topological assumptions imply their algebraic conditions. Wakker (1988, Thm 4.1) complements this by showing that, under the topological assumptions, any additive representation must be continuous. \square

8.5. Revealed trade-offs

Joint independence should not be confused with the weaker assumption of *single-factor independence* (also called *weak separability*): that preferences over each *single* factor do not depend on what is held fixed for the other factors. The fact that such independence must hold for preferences over multiple factors is the reason for the term “joint” in “joint independence”.

For example, any monotone preference relation on \mathbb{R}_+^J satisfies single-factor independence. Yet even with the additional assumptions of Theorem 3, not all monotone preferences have additive representations. To illustrate this, we need only find a continuous and monotone function that cannot be transformed monotonically into an additive function; $U(x_1, x_2) = \min\{2x_1 + x_2, x_1 + 2x_2\}$ will do.

With just two factors, single-factor and joint independence are the same. This is why Theorem 3 assumed that there were at least three essential factors. To handle the two-factor case, one needs an assumption that is more explicit about the trade-offs Anna may make between factors. We now review such an axiomatization and its extension to J factors.

Let $j \in \{1, \dots, J\}$, let $a_j, b_j, c_j, d_j \in X_j$, and let $x_{-j}, y_{-j} \in X_{-j}$. Suppose

$$(x_{-j}, a_j) \preceq (y_{-j}, b_j) \quad \text{but} \quad (x_{-j}, c_j) \succ (y_{-j}, d_j).$$

This reveals (one might say) that, keeping fixed the comparison x_{-j} versus y_{-j} on factors other than j , the strength of preference for c_j over d_j on factor j weakly exceeds the strength of preference for a_j over b_j on factor j . For nonadditive utility, a switch in such revealed strength of preference could easily occur. That is, there might also be $w_{-j}, z_{-j} \in X_{-j}$ such that

$$(w_{-j}, a_j) \succ (z_{-j}, b_j) \quad \text{but} \quad (w_{-j}, c_j) < (z_{-j}, d_j).$$

If there is such a reversal then coordinate j reveals contradictory trade-offs. The assumption that *no coordinate reveals contradictory trade-offs* is also called *triple cancellation* for the case of $J = 2$ (Krantz et al. 1971) and *generalized triple cancellation* for $J \geq 2$ (Wakker 1989).

One way to put the “no contradictory trade-offs” condition into words is roughly as follows. When it holds, Anna could say—

When I choose among consumption paths for the week, one of my considerations is what I will do on Friday night. I prefer to go to a great movie on Friday rather than read a book. I also prefer to go to a dance rather than play poker. But the first comparison is more likely to swing my decision than the second one.

—without any qualification about what is happening the rest of the week.

Our next theorem derives an additive representation from this condition.

THEOREM 4. *Assume $J \geq 2$ and that the following hold:*

1. X_j is a connected for all $j = 1, \dots, J$;
2. \succ is continuous;
3. *no coordinate reveals contradictory trade-offs.*

Then \succ has an additive representation that is continuous.

Proof. See Krantz et al. (1971, Sec. 6.2) for the two-factor case. See Wakker (1989, Thm. III.6.6) for the J -factor case. \square

9. Linear preferences

9.1. Linear utility

Suppose that X is a convex subset of \mathbb{R}^n . Then a subclass of utility functions that are additive across dimensions of the space are those that are linear.

Let $u \in \mathbb{R}^n$ be the vector representation of the linear utility function. That is, $U(x) = \sum_{i=1}^n u_i x_i = u \cdot x$, where $u \cdot x$ denotes the inner product. A typical indifference curve is

$$\{x \in X \mid u \cdot x = \bar{u}\}.$$

That is, the indifference curves are parallel hyperplanes (more specifically, the intersection of such parallel hyperplanes with X), and there is a direction U —perpendicular to the hyperplanes—that points in the direction of increasing utility.

These two conditions are also sufficient for the existence of a linear utility representation. We merely take the normal vector of one of the hyperplanes, pointing in the direction of increasing utility, and this defines a linear utility representation.

9.2. Linearity axiom

What axioms on the preference relation give us linear and parallel indifference curves? Let's start with just the linearity of a particular indifference curve. Roughly speaking, we need that, for any two points on an indifference curve, any line passing through those points also lies on the indifference curve—at least where it intersects X . We can state this condition as follows.

AXIOM L1. For all $0 < \alpha < 1$ and all $x, y \in X$,

$$\alpha x + (1 - \alpha)y \sim x \iff x \sim y. \quad (2)$$

Because Axiom L1 is an “if and only if” condition, it is stronger than convexity of the indifference curves and really does imply that they are linear subspaces. Take x and y on the same indifference curve and let $z \in X$ be any point on the line through x and y . If z lies between x and y , then we must have $z \sim x$ because of the \Leftarrow direction of condition (2). If instead z lies beyond y , then y is a convex combination of x and z and we have $z \sim x$ because of the \Rightarrow direction of condition (2).

The second condition we want is for the translation of an indifference curve to be another indifference curve, so that the indifference curves are parallel. We might state this as follows.

AXIOM L2. For all $x, y \in X$ and $z \in \mathbb{R}^n$ such that $x + z \in X$ and $y + z \in X$,

$$x + z \sim y + z \iff x \sim y.$$

However, there is a more succinct way to combine Axioms L1 and L2.

AXIOM L3. For all $0 < \alpha < 1$ and all $x, y, z \in X$,

$$\alpha x + (1 - \alpha)z \sim \alpha y + (1 - \alpha)z \iff x \sim y. \quad (3)$$

Axioms L1 and L2 are equivalent to Axiom L3. By letting $z = y$ in condition (3), one obtains condition (2). Figure 1 illustrates Axiom L3 for z that is not colinear with x and y . The axiom then implies that $x \sim y$ if and only if $x' \sim y'$ and $x'' \sim y''$. One can see why this axiom implies that the indifference curves are parallel.

However, Axiom L3 does not ensure that there is a common direction of increasing utility. To add this property, we can strengthen L3 by stating it in terms of \succsim rather than \sim . We thus obtain the final statement of the linearity axiom.




FIGURE 1. An illustration of Axiom L3 or Axiom L: Axiom L3 states that $x \sim y$ if and only if $x' \sim y'$ and $x'' \sim y''$; Axiom L is the same but for \succsim .

AXIOM L (Linearity). For all $0 < \alpha < 1$ and all $x, y, z \in X$,

$$\alpha x + (1 - \alpha)z \succsim \alpha y + (1 - \alpha)z \iff x \succsim y.$$

9.3. Continuity and the representation theorem

We have controlled for everything but the dimensionality of the indifference curves. It is fine for X to be one huge indifference curve; we can use the trivially linear utility function that maps all of X to 0. Otherwise, we need the dimensionality of the indifference curves to be one less than the dimensionality of X , so that each curve is the intersection of X and a hyperplane.

For example, lexicographic preferences on \mathbb{R}^2 satisfy the linearity axiom, and each indifference curve is indeed a linear subspace: a point. However, lexicographic preferences do not have a linear utility representation because the dimensionality of the indifference curves is too small.

A continuity assumption suffices to get the dimensionality correct. Consider, for example, an indifference curve through an alternative x that is neither the best nor the worst element of X . If preferences are continuous, then the set of points worse than x is open and the set of points better than x is open; hence their union cannot be a connected set. This means that the indifference curve through x must separate these two sets, which in turn requires that the indifference curve's dimension be one less than the dimension of X .

Given the linearity axiom, continuity is guaranteed by the following Archimedean axiom (which would be weaker than continuity in the absence of linearity).

AXIOM 3 (Archimedean). For all $x, y, z \in X$, if $x \succ y \succ z$ then there exist α, β in $(0, 1)$ such that

$$\alpha x + (1 - \alpha)z \succ y \succ \beta x + (1 - \beta)z.$$

We now have our representation theorem.

THEOREM 5. *If \succsim satisfies the linearity and Archimedean axioms, then \succsim has a linear utility representation.*

Proof. As we will see in Sections 13, this proposition is similar to the expected utility theorem of von Neumann and Morgenstern (1944). However, this form is due to Jensen (1967, Thm. 8). It can be generalized to infinite-dimensional vector spaces. The simplest way is to develop a theory for an abstract generalization of convex sets, called “mixture spaces”, as in Herstein and Milnor (1953). A comprehensive treatment can be found in Fishburn (1970, Chaps. 8 and 10). \square

9.4. Linearity implies joint independence

Suppose there are J factors and each X_j is a convex subset of Euclidean space. Then $X = X_1 \times \cdots \times X_J$ is also a convex set. According to Theorem 5, if \succsim satisfies the linearity and Archimedean axioms, then \succsim has a linear utility representation, which can be written

$$U(x_1, \dots, x_J) = \sum_{j=1}^J u_j \cdot x_j.$$

Since such a representation is additive across factors, \succsim satisfies joint independence. Thus, with such a factor structure, a corollary to Theorem 5 is that the linearity and Archimedean axioms imply joint independence.

However, the Archimedean axiom is not needed for this implication. A direct proof of the stronger result, stated in the next proposition, clarifies the link between linearity and joint independence.

PROPOSITION 3. *If \succsim satisfies linearity then \succsim satisfies joint independence.*

Proof. Let $\{K, L\}$ be a partition of $\{1, \dots, J\}$, let $a, b \in X_K$, and let $c, d \in X_L$. Assume $(a, c) \succsim (b, c)$. We show that $(a, d) \succsim (b, d)$.

Linearity and $(a, c) \succsim (b, c)$ imply

$$\frac{1}{2}(a, c) + \frac{1}{2}(a, d) \succsim \frac{1}{2}(b, c) + \frac{1}{2}(a, d). \quad (4)$$

By simple algebra, we can swap the terms a and b on the right hand side of (4), so that this equation becomes

$$\frac{1}{2}(a, c) + \frac{1}{2}(a, d) \succsim \frac{1}{2}(a, c) + \frac{1}{2}(b, d). \quad (5)$$

Since (a, c) is a common term on both sides of the preference in (5), we obtain from a second application of linearity that $(a, d) \succsim (b, d)$. \square

10. Cardinal uniformity across factors

10.1. Ordinal versus cardinal uniformity

Return to the setting of Section 8, in which there are J factors and $X = X_1 \times \cdots \times X_J$. There we found conditions under which Anna’s preferences would have an additive representation

$$U(x_1, \dots, x_J) = \sum_{j=1}^J u_j(x_j).$$

Suppose that each factor has the same set of possible alternatives: $X_j = \hat{X}$ for $j = 1, \dots, J$. It is then meaningful to consider the special case in which preferences over each factor are the same—i.e., in which u_j and u_k represent the same ordinal preferences for $j \neq k$.

To state this condition in terms of preferences rather than utility, define the preference relation on factor 1 (for example) as follows. Pick $x_{-1} \in \prod_{j=2}^J X_j$ and then define $x_1 \succsim_1 x'_1$ by $(x_1, x_{-1}) \succsim (x'_1, x_{-1})$; single-factor independence implies that this preference ordering on X_1 is the same no matter which x_{-1} we use to define it. Then the preferences \succsim satisfy *ordinal uniformity across factors* if \succsim_j is the same for all j .

However, ordinal uniformity adds little tractability beyond that of a general additive representation. It would be more interesting that the within-factor utility functions u_j were mere positive affine transformations of each other. Then we could find a common utility function $u: \hat{X} \rightarrow \mathbb{R}$ and weights $\delta_1, \dots, \delta_J > 0$ such that

$$U(x_1, \dots, x_J) = \sum_{j=1}^J \delta_j u(x_j).$$

Such a representation has *cardinal uniformity across factors*.

For example, suppose that the factors are different time periods. Then such a representation is what we mean by *time-independent preferences with discounting*. We might normalize the factors by dividing by δ_1 ; then the new δ_1 equals 1 and we interpret δ_j as the discount factor from period j to period 1.

Such time independence of the utility representation simplifies economic models and allows one to derive stronger results. What is the essence of cardinal uniformity that allows for this? *We can separate the relative importance of the factors from the within-factor preferences.* The weights $\{\delta_1, \dots, \delta_J\}$ provide a meaningful measure of the importance of the factors. For example, in a multi-agent intertemporal general equilibrium model, the assumption that all consumers have the same discount factors (i.e., the same impatience) is a powerful restriction even if their within-period utility functions u are heterogeneous. Without time-independent preferences, we would not even be able to formulate such an assumption.

Ordinal uniformity of preferences does not imply cardinal uniformity of the utility representation. Cardinal uniformity requires additional restrictions, which we explore in the rest of this section.

10.2. Ranking of factors

As the preceding intertemporal consumption example illustrates, one of the key features of cardinal uniformity is that there is a measure $\{\delta_1, \dots, \delta_J\}$ of the relative importance of the different factors. In terms of preferences, we can define an ordinal version of this condition as follows.

AXIOM 4 (Joint ranking of factors). Suppose that preferences satisfy ordinal uniformity and let \succsim^* be the common-across-factors ordering on \hat{X} . Let $x^1, x^2, x^3, x^4 \in \hat{X}$ be such that $x^1 \succ^* x^2$ and $x^3 \succ^* x^4$. Let $K, L \subset \{1, \dots, J\}$ and let, for example,

x_K^1 be the element of X_K that equals x^1 on each coordinate. Then

$$(x_K^1, x_{K^c}^2) \succ (x_L^1, x_{L^c}^2) \iff (x_K^3, x_{K^c}^4) \succ (x_L^3, x_{L^c}^4). \quad (6)$$

Joint ranking of factors is the same as Savage’s P4 axiom, which we will see in Section 11. It yields an unambiguous weak order over the subsets of factors when we interpret $(x_K^1, x_{K^c}^2) \succ (x_L^1, x_{L^c}^2)$ in equation (6) as meaning that the set of factors K is more important than the set of factors L .

Ordinal uniformity across factors has no implications for the trade-offs across factors, whereas joint ranking of factors does. For example, consider an intertemporal consumption example with a single good and monotone preferences. Ordinal uniformity holds merely because preferences are monotone. Joint ranking of factors implies additional restrictions, such as the following.

1. If Anna prefers the consumption path (100, 98, 98) over the consumption path (98, 100, 98), then she must also prefer (40, 30, 30) over (30, 40, 30). This captures the idea that Anna is impatient—not a normative assumption but one that is, at least, easy to interpret and empirically motivated.
2. The word “joint” (in “joint ranking of factors”) applies because K and L can be arbitrary subsets of factors. Thus, suppose that Anna prefers the consumption path (100, 98, 98, 100, 98) over (98, 100, 100, 98, 98), which reflects fairly complicated trade-offs between the periods. Then she must also prefer (40, 30, 30, 40, 30) over (30, 40, 40, 30, 30). This goes well beyond impatience.

However, joint ranking of factors (with also joint independence and ordinal uniformity) is not yet sufficient for the existence of a cardinally uniform representation unless the choice set is rich enough for the axioms to impose enough restrictions. The problem is analogous to the way joint independence is sufficient for an additive representation when there are three factors but not when there are only two. We explore three ways to resolve this problem.

10.3. Enriching the set of factors

One approach is to enrich the set of factors so that it is infinite and perfectly divisible.

The first steps are to extend the definition of an additive representation to an infinite set of factors and then extend Theorem 3 to axiomatize the existence of such a representation. For this we refer the reader to Wakker and Zank (1999, Thm. 12). Joint independence remains the main behavioral assumption for additivity. Some additional technical assumptions are needed to handle the nonatomicity of the set of factors.

For example, suppose the set of factors is an interval $[0, 1]$ of time. A consumption path is a function $\tilde{x}: [0, 1] \rightarrow \tilde{X}$. Joint independence and additional technical assumptions might yield an additive representation of the form

$$U(\tilde{x}) = \int_0^1 u(\tilde{x}(j), j) dj;$$

here, for each j , $u(\cdot, j): \hat{X} \rightarrow \mathbb{R}$ is the within-factor utility for factor j . In this example, integration is with respect to the Lebesgue measure, so (a) any set of factors of Lebesgue measure 0 is negligible, in the sense that changing a consumption plan on such a set does not change utility, and (b) the set of factors is nonatomic. Such features hold generally in the representation of Wakker and Zank (1999) and are important for what follows. Absent nonatomicity, the set of factors may not be rich enough for joint ranking of factors to imply cardinal uniformity. However, we repeat that these are essentially restrictions on the setting and not on the behavior of the decision maker.

We now show that, in this example, joint ranking of factors is sufficient to go from an additive representation to one that is also cardinally uniform. To keep the illustration simple, we assume that \hat{X} contains only the three elements a, b, c . The main step is the following lemma.

LEMMA 1. *Assume ordinal uniformity and joint ranking of factors, and assume that $a \succ^* b \succ^* c$. Then the ratio*

$$\frac{u(b, j) - u(c, j)}{u(a, j) - u(c, j)} \quad (7)$$

is the same for a.e. j .

Before proving this lemma, observe that it allows us to obtain a cardinally uniform representation. Define $v(a) = 1$ and $v(c) = 0$, and let $v(b)$ be the common value of the ratio (7). Define also $\pi(j) = u(a, j) - u(c, j)$. Then

$$V(\tilde{x}) \equiv \int_0^1 \pi(j) v(\tilde{x}(j)) dj$$

equals $U - \int_0^1 u(c, j) dj$. Thus, V is an additive representation with cardinal uniformity.

Proof of Lemma 1. We prove a contrapositive: Given ordinal uniformity and an additive representation, if the ratio (7) is not the same for almost every j then joint ranking of factors does not hold.

Suppose (7) varies on a set of positive measure. Then there exist $\rho > 0$ and disjoint sets $K, L \subset [0, 1]$ such that

$$\frac{u(b, j) - u(c, j)}{u(a, j) - u(c, j)} > \rho \text{ for } j \in K \quad \text{and} \quad \frac{u(b, j) - u(c, j)}{u(a, j) - u(c, j)} < \rho \text{ for } j \in L. \quad (8)$$

Given nonatomicity, we can find $A \subset K$ and $B \subset L$ of strictly positive measure such that

$$\int_A (u(a, j) - u(c, j)) dj = \int_B (u(a, j) - u(c, j)) dj. \quad (9)$$

Let \tilde{x}^1 be equal to a on A and to c elsewhere; let \tilde{x}^2 be equal to a on B and c elsewhere. Observe that

$$U(\tilde{x}^1) - U(\tilde{x}^2) = \int_A (u(a, j) - u(c, j)) dj - \int_B (u(a, j) - u(c, j)) dj = 0. \quad (10)$$

That is, $\tilde{x}^1 \sim \tilde{x}^2$.

Joint ranking of factors implies that this same indifference should hold if \tilde{y}^1 and \tilde{y}^2 are defined accordingly except with a replaced by b . However, the two inequalities

$$\begin{aligned} \int_A (u(b, j) - u(c, j)) dj &> \rho \int_A (u(a, j) - u(c, j)) dj, \\ \int_B (u(b, j) - u(c, j)) dj &< \rho \int_B (u(a, j) - u(c, j)) dj \end{aligned}$$

(which following from the two inequalities in equation (8)), together with equality (9), imply that

$$\int_A (u(b, j) - u(c, j)) dj > \int_B (u(b, j) - u(c, j)) dj.$$

Therefore, as an analog to equation (10), we have $U(\tilde{y}^1) > U(\tilde{y}^2)$. Hence, joint ranking of factors does not hold. \square

10.4. A factor-independent version of no contradictory trade-offs

The second approach is to assume that \hat{X} is connected and then use, as a starting point, the axiom that no coordinate reveals contradictory trade-offs. We capture state-independence by modifying the axiom so that it holds even when we permute the indices. That is, if

$$(x_{-j}, a) \preceq (y_{-j}, b) \quad \text{but} \quad (x_{-j}, c) \succ (y_{-j}, d),$$

then it cannot be the case that

$$(w_{-k}, a) \succ (z_{-k}, b) \quad \text{but} \quad (w_{-k}, c) < (z_{-k}, d).$$

(We are being loose in the notation: the idea is that a, b, c, d show up in coordinate j in the first equation but in coordinate k in the second equation.) The resulting axiom is called *cardinal coordinate independence* or *no contradictory trade-offs* in Wakker (1989).

PROPOSITION 4 (Wakker 1989). *Suppose that \hat{X} is connected and that all factors are essential. Assume \succ is continuous and satisfies cardinal coordinate independence. Then \succ has a continuous additive representation that is cardinally uniform across factors.*

Cardinal coordinate independence subsumes joint independence, ordinal uniformity across factors, and joint ranking of factors. It captures the idea that the relative strength of trade-offs is uniform across factors.

10.5. The linear case

The third approach is to consider linear preferences. Suppose that \hat{X} is a convex subset of Euclidean space and that we obtain an additive representation in

which each u_j is linear. Recall that, if u_1 and u_2 (for example) are two linear representations of the same ordinal preferences on \hat{X} , then u_2 is a positive affine transformation of u_1 . Thus, linearity combined with the ordinal uniformity across factors gives us cardinal uniformity across factors!

We summarize this in a single representation theorem. If \hat{X} is convex, then $X = \hat{X}^J$ is also convex and so we can impose the linearity and Archimedean axioms directly on \succsim . Linearity implies additivity; the axiom of joint independence is thus redundant. We merely need to add ordinal uniformity.

PROPOSITION 5. *Suppose \succsim satisfies the linearity, Archimedean, and ordinal uniformity across factors axioms. Then \succsim has a utility representation of the form*

$$U(x_1, \dots, x_J) = \sum_{j=1}^J \delta_j u(x_j),$$

where $u: \hat{X} \rightarrow \mathbb{R}$ is linear and continuous.

Proof. This is a consequence of Theorem 5 and the observation that ordinal uniformity implies cardinal uniformity for a linear representation. In Section 15, we will see that this is the expected utility theorem of Anscombe and Aumann (1963). \square

11. States of nature and acts

We now reach (finally!) decisions under uncertainty. The first question is basic: How do we model uncertainty? We start with the states-of-the-world formulation and then move to the more reduced-form lotteries formulation.

11.1. Setup

In a static decision problem under uncertainty, Anna is uncertain about the outcomes of her actions. A formal model should make a careful distinction between those aspects of the world:

1. that she controls, which we call an *action* and denote by $a \in A$;
2. that she is uncertain about and cannot influence, which we call a *state* and denote by $s \in S$;
3. that she cares directly about, which we call an *outcome* and denote by $z \in Z$.

These aspects are linked by a function $F: A \times S \rightarrow Z$ that summarizes how the outcome is determined by her actions (which she controls) and the state (which she cannot influence).

11.2. From large to small worlds

Savage (1954, p. 9) describes a state of the world as “a description of the world, leaving no relevant aspect undescribed”, but in Chapter 5 he goes on to explain

that an actual model of decisions (whether constructed by a decision theorist or contemplated by Anna) could not contain such detail. He then draws the following link between the abstract “large worlds”, whose specification can omit no details, and the “small worlds” that appear in our models.

Let S^* be the set of “large worlds” or states as envisioned by Savage. A subset $E \subset S^*$ is called an *event*. To construct a model of a particular decision problem, we replace S^* by a partition $\{E_1, \dots, E_n\}$ of S^* that is

- *fine* enough that, for every action, states in the same event lead to the same outcome,
- but otherwise as *coarse* as possible.

Each state s_j in the model corresponds to an event E_j in the partition.

Consider, for example, a portfolio choice problem in which Anna will invest a fixed amount of money in securities, hold the securities for one year, and then cash out the securities. The actions are the different portfolios Anna could choose. Suppose that Anna cares only about the monetary payoff of her portfolio; such payoffs are the outcomes. Then a state could represent the values of all the securities at the time Anna cashes in her portfolio. Each “state” in the model is really an event. For instance, distinct large-world states within an event could differ in terms of the outcome of elections, the growth of the plants in Anna’s garden, and other factors that do not affect portfolio returns.

11.3. From actions to acts

The decision maker chooses an action $a \in A$. For example, Anna chooses a portfolio. Given action a , each state s leads to outcome $F(a, s)$. This defines a map $s \mapsto F(a, s)$ from S to Z . We call such a map an *act*. The act $s \mapsto F(a, s)$ is the one induced by the action a .

A first assumption is that, when choosing among actions, Anna compares the acts that they induce. This is plausible as long as we have defined outcomes carefully enough that they include any aspects of the action that Anna may care about (see Section 4). This leads to a reduced-form model in which Anna chooses among acts. Let’s go a step further. Assume that Anna can contemplate hypothetical choices among acts, even those that are not induced by any action.

(Our distinction between actions and acts abuses vocabulary but is useful. When Savage coined the term “act”, he was thinking of it as synonymous with “action”. However, he jumped immediately to the reduced form and defined an act to be a mapping between S and Z . It is useful to retain the notion of the true action that Anna controls so that the model can more naturally describe actual decision problems.)

We can relate the resulting model to the ones developed earlier in the chapter as follows.

1. The set X of alternatives from which Anna chooses is the set of acts.
2. This set has a product structure in which the set of factors is the set S of states. Whereas we denoted an alternative by $x = (x_1, \dots, x_J)$ in the factor model, here we use the equivalent functional notation $f: S \rightarrow Z$.




FIGURE 2. Visualization of an act.

3. In each state, the set of outcomes is the same: Z . In the factor decomposition, this corresponds to the special case in which $X_j = \hat{X} = Z$ for each factor.

Although mathematically an act is just an element of a product set, it is appealing to visualize an act (when there are finitely many states) as in Figure 2, because it illustrates the unfolding of uncertainty. At the moment of making a choice, Anna stands in a position at or before the root node. Then one of the three states is realized—that is, one and only one of the branches from the root node occurs, leading to the consumption bundle (or, more generally, the outcome) $z = f(s)$.

12. Sure-thing principle and additivity

Since X has a product structure, we can look for an additive representation of Anna's preferences \succsim .

12.1. Joint independence

In this setting, the assumption of joint independence was called the *sure-thing principle* by Savage and has also been called the *independence axiom* (although the latter term has been used for related axioms in other settings). We pointed out that joint independence was a strong assumption when the factors are consumption of different goods or consumption in different time periods: one would never insinuate in those settings that joint independence was a normative axiom as opposed to a mere simplifying assumption. However, in this decision problem under uncertainty, the sure-thing principle is a plausible normative assumption.

With state-contingent consumption, only one of the states will ever be realized. It seems natural to most people, barring any mistakes Anna might make, that deciding how to spend the money Anna might earn if she wins a bet does not depend on what she might do if she loses the bet. We argue this more fully in the following section.

12.2. Dynamic consistency

The notion of dynamic consistency can help us understand the sure-thing principle and judge whether it is a compelling normative axiom. Because our goal is not to develop a full theory of dynamic consistency for its own sake and because such a framework is notationally intensive when set out formally, we keep the

discussion informal.

In Section 3, we stated that (a) our model of static choice also applied to choosing a once-and-for-all plan in a temporal setting, but that (b) this chapter was not about making decisions at different decision epochs or the link between such decisions. Consider, however, such a possibility.

We are developing a normative model. Hence, though it would difficult for her to do so, imagine that Anna can envision all the opportunities in which she will make decisions, that any uncertainty about available actions and outcomes of actions is reflected in her model of a set of states of the world, and that Anna understands how uncertainty is resolved over time as she learns new information. (All this would be modeled formally as a decision tree, which means an extensive game form in which there are moves by nature but otherwise Anna is the only player.) A plan specifies the actions Anna would choose at each decision epoch (i.e., a plan is the same as a strategy in an extensive game form). Anna can then envision all possible plans and choose among them.

Suppose she then begins executing her preferred plan. At each decision epoch, she could reevaluate the plan—that is, consider all continuation plans on the subtree (subgame) that follows and choose among them. Dynamic consistency means that, as long as she has followed the ex ante preferred plan so far, Anna never wishes to deviate from the plan.


This paints an unrealistic picture of the way people actually make decisions. We do not set up a full model and choose among all possible plans. Even in an easily codified situation such as a chess game to be played during an afternoon, the set of possibilities is so large that it overwhelms our ability to envision and choose among them. Thus, we always have partially formulated plans, we constantly extend our partial model further into time, we review our available options, and we revise our decisions.

This is all due to the sheer complexity of devising a complete dynamic plan. However, it is fairly well accepted that dynamic consistency is a reasonable normative axiom—something Anna would satisfy if she were infinitely smart and had unlimited computational power.

Consider a simple situation in which there are six states, $\{s_1, s_2, s_3, s_4, s_5, s_6\}$, and in which Anna learns whether the state is in $E_1 \equiv \{s_1, s_2, s_3\}$ or in $E_2 \equiv \{s_4, s_5, s_6\}$ before choosing an action. She formulates a plan of what action to choose depending on whether she observes E_1 or E_2 . Each such plan leads to an act.

For a given act f , we can represent the unfolding of information as in Figure 3. Observe that this is not a full decision tree; it depicts a single act and does not show the actions available at each decision epoch. Only nature's moves from the decision tree are shown. Behind the scenes we understand that different decisions lead to different acts, which means that Anna's choices determine which outcomes are at the leaves. The partial acts f_1 and f_2 specify the outcomes for the states in E_1 and E_2 , respectively. We can write $f = (f_1, f_2)$.

If Anna is dynamically consistent, then if she observed E_1 she will want to stick to her plan. In other words, when choosing among acts defined on the states in E_1 after observing that event, she will still choose f_1 among the available acts. It is a fairly compelling assumption that, when she chooses among continuation

FIGURE 3. An act that results from making decisions after observing $\{E_1, E_2\}$.

plans having observed E_1 , her preferences do not depend on what she thinks she might have done had E_2 been realized instead.

No theorem states that an infinitely smart person capable of such complex planning could not then be influenced in her decisions at each decision epoch by the plans she had made if other contingencies had been realized. “Fairly compelling assumption” simply means that these authors (at least) imagine that if we were so smart then this is how we would make decisions.

This assumption, combined with dynamic consistency, implies that if $(f_1, f_2) \succcurlyeq (f'_1, f'_2)$ then also $(f_1, f'_2) \succcurlyeq (f'_1, f'_2)$ for any other partial act f'_2 defined on the states in E_2 . More generally, it implies that preferences over acts satisfy the sure-thing principle.

12.3. Expected utility representation with state-dependent utility

Assume now that the sure-thing principle holds (along with other technical assumptions, such as those in Theorem 3), so that Anna’s preferences over acts have an additive representation $\sum_{j=1}^J u_j(x_j)$, which we may write as

$$U(f) = \sum_{s \in S} u_s(f(s)).$$

Let $\pi: S \rightarrow \mathbb{R}$ be any probability measure on S such that $\pi(s) > 0$ for all $s \in S$. For each $s \in S$, define $v_s: Z \rightarrow \mathbb{R}$ by $v_s(z) = u_s(z)/\pi(s)$. Then

$$U(f) = \sum_{s \in S} \pi(s) v_s(f(s)). \quad (11)$$

Written this way, it is tempting to interpret U as an expected utility representation. Anna’s utility in state s is $v_s(f(s))$ and, given the probability measure π on states, $U(f)$ is the expected value of Anna’s utility. However, since the probability measure was chosen arbitrarily, the probabilities or beliefs used in this representation are indeterminate. Therefore, we cannot claim to have uncovered Anna’s beliefs from her preferences.

The problem arises because the utility functions u_s in the additive representation have no link to each other. Utility is therefore said to be *state dependent*.

12.4. Expected utility representation with state-independent utility

On the other hand, suppose that the representation satisfies cardinal uniformity across states and can thus be written as

$$U(f) = \sum_{s \in S} \delta_s u(f(s)),$$

where $\delta_s > 0$ for all s and $u: Z \rightarrow \mathbb{R}$. We then say that the utility is *state independent*. We can normalize the weights $\{\delta_s \mid s \in S\}$ so that they sum to 1 by setting $\pi(s) = \delta_s / (\sum_{s' \in S} \delta_{s'})$; this is equivalent to multiplying the utility representation by $1 / (\sum_{s' \in S} \delta_{s'})$. Then we have the representation

$$V(f) = \sum_{s \in S} \pi(s) u(f(s)). \quad (12)$$

If we take another additive representation, we know that it is a positive affine transformation of V . When we retain the restriction that the weights π sum to 1, the multiplicative and additive constants must end up in the function u . Thus, any additive representation of the preferences can be written in the form of (12) and will always end up with the same weighting function π . It then becomes more plausible to refer to these weights as probabilities or beliefs because they are uniquely identified.

12.5. Is identification of beliefs important?

The answer is both “no” and “yes”, as we now explain.

Identification is a leap of faith. Although the weights π in representation (12) are uniquely defined, their interpretation as beliefs remains a leap of faith. We cannot fully disentangle—from preferences over acts—whether Anna cares about what happens in a state because she considers the state likely or because the state affects her feelings. This is a general point, but we can give a concrete technological example in which preferences have the representation in equation (12) yet the weights $\pi(s)$ do not represent beliefs. Let the elements of Z be inputs in a single-output production process that is affected by state-contingent shocks such that output in state s , given inputs $z \in Z$, is $\theta_s u(z)$ (i.e., the production function in state s is $z \mapsto \theta_s u(z)$). An act is a state-contingent input plan $f: S \rightarrow Z$. Anna’s beliefs are given by $\hat{\pi}$, and she is a risk-neutral expected utility maximizer with respect to output. Then her preferences over state-contingent input plans are given by

$$V(f) = \sum_{s \in S} \hat{\pi}(s) \theta_s u(f(s)).$$

The weights $\hat{\pi}(s) \theta_s$, written as $\pi(s)$ in equation (12), confound probabilities and productivity shocks. This is an argument that Karni (1985) gives in favor of studying state-dependent preferences and utility as opposed to inferring beliefs from a normalization of weights that is not derived from preferences or choice behavior.

Identification of beliefs is not needed for Bayesian decision making. Are we concerned merely that Anna act as if she were probabilistically sophisticated and maximized expected utility, so that we can apply the machinery of Bayesian statistics to Anna's dynamic decision making? Or rather is our objective to uniquely identify her beliefs?

The latter might be useful if we wanted to measure beliefs from empirically observed choices in one decision problem in order to draw conclusions about how Anna would act in another decision problem. Otherwise, the former is typically all we need, and state-dependent preferences are sufficient.

We can pick an additive representation of the form (11) with any weights π . Suppose that Anna faces a dynamic decision problem in which she can revise her choices at various decision nodes after learning some information (represented by a partition of the set of states). Given dynamic consistency, she will make the same decisions whether she makes a plan that she must adhere to or instead revises her decisions conditional on her information at each decision node. Furthermore, in the latter case her preferences over continuation plans will be given by expected utility maximization with the same state-dependent utilities and with weights (beliefs) that are revised by Bayesian updating. This may allow the analyst to solve her problem by backward induction (dynamic programming or recursion), thereby decomposing a complicated optimization problem into multiple simpler problems.

Yet state independence is a powerful restriction. The real power of state-independent utility comes from the structure and restrictions that this imposes on preferences, particularly in equilibrium models with multiple decision makers. We already discussed this in the context of an intertemporal model with cardinally uniform utility. Let's revisit this point in the context of decision making under uncertainty.

With state-independent utility, we can separate the relative probabilities of the states from the preferences over outcomes. For example, when the outcomes are money, we can separate beliefs from risk preferences. This is particularly powerful in a multi-person model because we can then give substance to the assumption that all decision makers have the same beliefs. Consider a general equilibrium model of trade in state-contingent transactions, such as insurance or financial securities. Suppose that all traders have state-independent utility with the same beliefs but heterogeneous utilities over money. If the traders' utilities are strictly concave (they are risk averse) and if the total amount of the good that is available is state independent (no aggregate uncertainty), then in any Pareto efficient allocation each trader's consumption is state independent (each trader bears no risk).

State independence is without loss of generality (more or less). It can be argued that state independence is without loss of generality: if it is violated, one can redefine outcomes to ensure that the description of an outcome includes everything Anna cares about—even things that are part of the description of the state. However, when this is done some acts are clearly hypothetical.

Perhaps the two states are "Anna's son has a heart attack" and "Anna's son heart is just fine". What Anna controls is whether her son has heart surgery.

Clearly her preferences for heart surgery depend on whether or not her son has a heart attack. However, we can define an outcome so that it is specified both by whether her son has a heart attack and by whether he undergoes the surgery. In order to maintain the assumption that the set of acts is the set of all functions from states to outcomes, Anna must be able to contemplate and express preferences among such hypothetical acts as the one in which her son has a heart attack and gets heart surgery in both states, including the state in which he does not have a heart attack!

Furthermore, when decision under uncertainty is applied to risk and risk sharing, the modeler assumes that preferences over *money* are state independent. This is a strong assumption even if preferences were state independent for some appropriately re-defined set of outcomes.

13. Lotteries

13.1. From subjective to objective uncertainty

We postpone until Section 14 a discussion of the axioms that capture state independence of preferences and yield a state-independent representation $U(f) = \sum_{s \in S} \pi(s) u(f(s))$. In the meantime, we consider how state independence combined with objective uncertainty allows for a reduced-form model in which choices among state-dependent outcomes (acts) is reduced to choices among probability measures on outcomes (lotteries). We then axiomatize expected utility for such a model.

One implication of state-independent expected utility is that preferences depend only on the probability measures over outcomes that are induced by the acts. That is, think of an act f as a random object whose distribution is the induced probability measure p on Z . Assume that S and Z are finite, so that this distribution is defined by $p(z) = \pi\{s \in S \mid f(s) = z\}$. We can then rewrite $U(f) = \sum_{s \in S} \pi(s) u(f(s))$ as $\sum_{z \in Z} p(z) u(z)$. In particular, Anna is indifferent between any two acts that have the same induced distribution over outcomes.

Let us now take as our *starting point* that Anna's decision problem reduces to choosing among probability measures over outcomes—without a presumption of having identified an expected utility representation in the full model. We then state axioms within this reduced form that lead to an expected utility representation.

For this to be an empirical exercise (i.e., in order to be able to elicit preferences or test the theory), the probability measures over outcomes must be observable. This means that the probabilities are generated in an objective way, such as by flipping a coin or spinning a roulette wheel. Therefore, this model is typically referred to as one of objective uncertainty. The other reason to think of this as a model of objective uncertainty is that we will need data on how the decision maker would rank all possible distributions over Z . This is plausible only if we can generate probabilities using randomization devices.

Thus, the set of alternatives in Anna's choice problem is the set of probability measures defined over the set Z of outcomes. To avoid the mathematics of




FIGURE 4. A lottery.

measure theory and abstract probability theory, we continue to assume that Z is finite, letting n be the number of elements. We call each probability measure on Z a *lottery*.

Let P be the set of lotteries. Each lottery corresponds to a function $p: Z \rightarrow [0, 1]$ such that $\sum_{z \in Z} p(z) = 1$. Each $p \in P$ can equivalently be identified with the vector in \mathbb{R}^n of probabilities of the n outcomes. The set P is called the *simplex* in \mathbb{R}^n ; it is a compact convex set with $n - 1$ dimensions.

We can illustrate a lottery graphically as in Figure 4. The leaves correspond to the possible outcomes and the edges show the probability of each outcome. Figure 4 looks similar to the illustration of an act in Figure 2, but the two figures should not be confused. When Anna considers different acts, the states remain fixed in Figure 2 (as do their probabilities); what change are the outcomes. When Anna considers different lotteries, the outcomes remain fixed in Figure 4; what changes are the probabilities. This reduced-form model of lotteries has a flexibility with respect to possible probability measures over outcomes that would not be possible in the states model unless the set of states were uncountably infinite and beliefs were nonatomic.

By an *expected utility* representation of Anna's preferences \succsim on P we mean one of the form

$$U(p) = \sum_{z \in Z} p(z) u(z),$$

where $u: Z \rightarrow \mathbb{R}$. Then $U(p)$ is the expected value of u given the probability measure p on Z . We call this a *Bernoulli representation* because Bernoulli (1738) posited such an expected utility as a resolution to the St. Petersburg paradox: that a decision maker would prefer a finite amount of money to a gamble whose expected payoff was infinite. Bernoulli took the utility function $u: Z \rightarrow \mathbb{R}$ as a primitive and expected utility maximization as a hypothesis. His innovation was to allow for an arbitrary, even bounded, function $u: Z \rightarrow \mathbb{R}$ for lotteries over money rather than a linear function, thereby avoiding the straitjacket of expected value maximization—the state of the art in his day.

Expected utility did not receive much further attention until von Neumann and Morgenstern (1944) first axiomatized it (for use with mixed strategies in game theory). For this reason, the representation is also called a von Neumann–Morgenstern utility function. As we do here, von Neumann and Morgenstern took preferences over lotteries as a primitive and uncovered the expected utility representation from several axioms on those preferences.

13.2. Linearity of preferences

Recall that P is a convex set and recall from Section 9 that \succsim admits a linear utility representation if it satisfies the linearity and Archimedean axioms. We proceed as follows.

1. We observe that a linear utility representation is the same as a Bernoulli representation.
2. We discuss the interpretation of the linearity and Archimedean axioms.

In this setting, linearity (Axiom L) is called the *independence axiom*.

So suppose we have a linear utility representation

$$U(p) = \sum_{z \in Z} u_z p_z$$

of Anna's preferences. We can write the vector $\{u_z \mid z \in Z\}$ of coefficients as a function $u: Z \rightarrow \mathbb{R}$ and use the functional form $p: Z \rightarrow [0, 1]$ of a lottery p . Then the linear utility representation can be written as

$$U(p) = \sum_{z \in Z} p(z) u(z), \quad (13)$$

that is, as a *Bernoulli representation*. Like any additive representation, this one is unique up to a positive affine transformation; such a transformation of U corresponds to an affine transformation of u . All this is summarized in our next theorem.

THEOREM 6. *If \succsim satisfies the linearity (independence) and Archimedean axioms, then \succsim has a Bernoulli representation.*

Proof. This is an application of Theorem 5; as such, it is due to Jensen (1967, Thm. 8). Von Neumann and Morgenstern's representation theorem used a different set of axioms that implied but did not contain an explicit independence (linearity) axiom like our Axiom L. The role of the independence axiom, which we interpret further in what follows, was uncovered gradually by subsequent authors. See Fishburn and Wakker (1995) for a history of this development. \square

13.3. Interpretation of the axioms

The convex combinations that appear in the linearity and Archimedean axioms have a nice interpretation in our lotteries setting. Suppose the uncertainty by which outcomes are selected unfolds in two stages. In a first stage, there is a random draw to determine which lottery is faced in a second stage. With probability α , Anna faces lottery p in the second stage; with probability $1 - \alpha$ she faces lottery r . This is called a *compound lottery* and is illustrated in Figure 5.

Consider the overall lottery t that Anna faces ex ante, before any uncertainty unfolds. The probability of outcome z_1 (for example) is $t_1 = \alpha p_1 + (1 - \alpha)r_1$. As a vector, the lottery t is the convex combination $\alpha p + (1 - \alpha)r$ of p and r . Thus, we can interpret convex combinations of lotteries as compound lotteries.




FIGURE 5. A compound lottery.




FIGURE 6. Another compound lottery.

Consider this compound lottery and the one in Figure 6, recalling the discussion of dynamic consistency and the sure-thing principle from Section 12. Suppose Anna chooses t over t' and then, after learning that she faces lottery p in the second stage, is allowed to change her mind and choose lottery q instead. Dynamic consistency implies that she would not want to do so. Furthermore, analogous to our normative justification of the sure-thing principle, it is also natural that her choice between p and q at this stage would depend neither on which lottery she would otherwise have faced along the right branch of the first stage nor on the probability with which the left branch was chosen. Together, these two observations imply that she would choose lottery t over t' if and only if she would choose lottery p over q . Mathematically, in terms of the preference ordering \succsim , this is Axiom L. It is called the *independence axiom* or *substitution axiom* in this setting because the choices between t and t' are then independent of which lottery we substitute for r in Figure 6.

Thus, the justification for the independence (linearity) axiom in this lotteries model is the same as for the sure-thing principle (joint independence axiom) in the states model, but the two axioms are mathematically distinct because the two models define the objects of choice differently (lotteries vs. acts).

The Archimedean axiom has the following meaning. Suppose that Anna prefers lottery p over lottery q . Now consider the compound lottery t in Figure 6. Lottery r might be truly horrible. However, if the Archimedean axiom is satisfied then, as long as the right branch of t occurs with sufficiently low probability, Anna still prefers lottery t over lottery q . This is illustrated by the risk of

death that we all willingly choose throughout our lives. Death is certainly something “truly horrible”; however, every time we cross the street we choose a lottery with small probability of death over the lottery we would face by remaining on the other side of the street.

13.4. Calibration of utilities

The objective probabilities are used in this representation to calibrate the decision maker’s strength of preference over the outcomes. To illustrate how this is done, suppose Anna is considering various alternatives that lead to varying objectively measurable probabilities of the following outcomes:

- e — Anna stays in her current employment;
- m — Anna gets an MBA but then does not find a better job;
- M — Anna gets an MBA and then finds a much better job.

We let $Z = \{e, m, M\}$ be the set of outcomes and, since this is a reduced form, we view her choice among her actions as boiling down to the choice among the probabilities over Z that the actions induce. Furthermore, we suppose that she can contemplate choices among all probability measures on Z , and not merely those induced by one of her actions. We assume $M \succ e \succ m$, where (for example) $M \succ e$ means that she prefers getting M for sure to getting e for sure.

Anna’s preference for e relative to m and M can be quantified as follows. We first set $u(M) = 1$ and $u(m) = 0$. We then let $u(e)$ be the unique probability for which she is indifferent between getting e for sure and the lottery that yields M with probability $u(e)$ and m with probability $1 - u(e)$ —that is, for which $e \sim u(e)M + (1 - u(e))m$. The closer e is to M than to m in her strength of preference, the greater this probability $u(e)$ would have to be and, in our representation, the greater is the utility $u(e)$ of e .

The Archimedean axiom implies that such a probability $u(e)$ exists. The independence axiom then implies that the utility function $u: Z \rightarrow \mathbb{R}$ thus defined yields a Bernoulli representation of Anna’s preferences. The actual proof of the representation theorem is an extension of this constructive proof to more general Z .

14. Subjective expected utility without objective probabilities

14.1. Overview

Let us return to the states and acts setting of Sections 11 and the state-dependent expected utility representation from Section 12. Recall the challenge—posed but not resolved—of finding a *state-independent* representation, so that the probabilities would be uniquely identified and could be interpreted as beliefs revealed by the preferences over acts. This is called *subjective expected utility* (SEU).

One of the first derivations of subjective expected utility (involving the joint derivation of subjective probabilities to represent beliefs about the likelihood of events as well as the utility index over outcomes) appeared in a 1926 paper by

Frank Ramsey, the English mathematician and philosopher. This article was published posthumously in Ramsey (1931) about the same time that an independent but related derivation appeared in Italian by the statistician de Finetti (1931). The definitive axiomatization in a purely subjective uncertainty setting appeared in Leonard Savage's 1954 book, *The Foundation of Statistics*.

In Section 12, we showed that the sure-thing principle implied additivity of the utility. We went on to say that SEU requires that the additive utility be cardinally uniform across states, but we stopped before showing how to obtain such a conclusion. Recall further back to Section 10, where we tackled cardinal uniformity in the abstract factors setting. Axiomatizing cardinal uniformity was tricky, but we outlined three solutions. Each of those solutions corresponds to an approach taken in the literature on subjective expected utility.

1. Savage (1954) used an infinite and nonatomic state space as in Section 10.3. We develop this further in Section 14.2.
2. Wakker (1989) assumed a connected (hence infinite) set of outcomes and assumed cardinal coordinate independence, as we did in Section 10.4. Cardinal coordinate independence involves specific statements about how the decision maker treats trade-offs across different states and assumes that such trade-offs are state independent.
3. Anscombe and Aumann (1963) mixed subjective and objective uncertainty to obtain a linear representation, as in Section 10.5. We develop this in Section 15.

14.2. Savage

We give a heuristic presentation of the representation in Savage (1954). (In what follows, P_n refers to Savage's numbering of his axioms.) Savage began by assuming that preferences are transitive and complete (P1: weak order) and satisfy joint independence (P2: sure-thing principle); this yields an additive or state-dependent representation. The substantive axioms that capture state independence are ordinal uniformity (P3: ordinal state independence) and joint ranking of factors (P4: qualitative probability).

As a normative axiom, P3 is really a statement about the ability of the modeler to define the set of outcomes so that they encompass everything that Anna cares about. Then, given any realization of the state, Anna's preferences over outcomes should be the same.

Because Savage works with an infinite state space in which any particular state is negligible, his version of P3 is a little different from ours and he needs an additional related assumption. These are minor technical differences.

1. Savage's P3 states that Anna's preferences are the same conditional *on any nonnegligible event*, rather than on any state. With finitely many states, the two axioms are equivalent.
2. Savage adds an axiom (P7) that the preferences respect statewise dominance: given Anna's state-independent ordering \succ^* on Z , if f and g are such that $f(s) \succ^* g(s)$ for all $s \in S$ then $f \succ g$. With finitely-many states, this condition is implied by the sure-thing principle and ordinal state independence.

Let consider in more detail Savage's P4, which is our joint ranking of factors. We begin by restating this axiom using the terminology and notation of the preferences-over-acts setting.

AXIOM 5 (Qualitative probability). Suppose that preferences satisfy ordinal state independence and let \succ^* be the common-across-states ordering on Z . Let $A, B \subset S$ be two events. Suppose that $z^1 \succ^* z^2$ and $z^3 \succ^* z^4$. Let, for example, $(I_A z^1, I_{A^c} z^2)$ be the act that equals z^1 on event A and z^2 on its complement. Then

$$(I_A z^1, I_{A^c} z^2) \succ (I_B z^1, I_{B^c} z^2) \iff (I_A z^3, I_{A^c} z^4) \succ (I_B z^3, I_{B^c} z^4).$$

This axiom takes state independence one step further: it captures the idea that the decision maker cares about the states only because they determine the likelihood of the various outcomes determined by acts. If preferences are state independent, then the only reason Anna would prefer $(I_A z^1, I_{A^c} z^2)$ to $(I_B z^1, I_{B^c} z^2)$ is because she considers event A to be more likely than event B . In such case, she must also prefer $(I_A z^3, I_{A^c} z^4)$ to $(I_B z^3, I_{B^c} z^4)$.

As explained in Section 10.3, ordinal state independence and qualitative probability impose enough restrictions to yield state-independent utility only if the choice set is rich enough—with one approach being to have a nonatomic set of factors or states. This is the substance of Savage's axiom P6 (continuity). The richer state space allows one to calibrate beliefs separately from payoffs over the outcomes.

15. Subjective expected utility with objective probabilities

15.1. Horse-race/roulette-wheel lotteries

Anscombe and Aumann (1963) avoided resorting to an infinite state space or axioms beyond joint independence and ordinal uniformity by combining (a) a lotteries framework with objective uncertainty and (b) a states framework with subjective uncertainty.

In their model, an act assigns to each state a lottery with objective probabilities. These two-stage acts are also called horse-race/roulette-wheel lotteries, but we continue to refer to them merely as acts and to the second-stage objective uncertainty as lotteries. Fix a finite set S of states and a finite set Z of outcomes. We let P be the set of lotteries on Z . An act is a function $f: S \rightarrow P$. Let H be the set of acts.

15.2. Linearity: sure-thing principle and independence axiom

First notice that H , which is the product set P^S , is also a convex set and that the convex combination of two acts can be interpreted as imposing compound lotteries in the second (objective) stage of the unfolding of uncertainty. In other words, for any pair of acts f, g in H and any α in $[0, 1]$, $\alpha f + (1 - \alpha)g$ corresponds to the act h in H for which $h(s) = \alpha f(s) + (1 - \alpha)g(s)$, where $\alpha f(s) + (1 - \alpha)g(s)$ is the convex combination of lotteries $f(s)$ and $g(s)$.

In Section 9, we showed that \succsim has a linear utility representation if \succsim satisfies the linearity and Archimedean axioms. Let us consider the interpretation of such a utility representation and the interpretation of these axioms.

The dimensions of H are $S \times Z$ and a linear utility function on H can be written as

$$\sum_{s \in S} \sum_{z \in Z} u_{sz} p_{sz} = \sum_{s \in S} \sum_{z \in Z} u_{sz} \times f(s)(z). \quad (14)$$

On the left side, we have represented the element of H as a vector $p \in \mathbb{R}^{S \times Z}$; the probability of outcome z in state s is p_{sz} . On the right side, we have represented the element of H as an act $f: S \rightarrow P$; the probability of outcome z in state s is $f(s)(z)$. (That is, $f(s)$ is the probability measure or lottery in state s and $f(s)(z)$ is the probability assigned to z by that measure.) We used a “ \times ” on the right-hand side for simple multiplication to make clear that $f(s)(z)$ is a single scalar term. The order of summation in equation (14) is irrelevant.

For any probability measure π on S we can also write the linear utility function in the form

$$\sum_{s \in S} \pi(s) \sum_{z \in Z} f(s)(z) \times u_s(z). \quad (15)$$

We derived equation (15) from (14) by:

- dividing each coefficient u_{sz} by $\pi(s)$ and writing the result as $u_s(z)$; then
- reversing the order of multiplication so that $\sum_{z \in Z} f(s)(z) \times u_s(z)$ is recognized as the expected utility in state s —given that $f(s)$ is the probability measure on Z and $u_s: Z \rightarrow \mathbb{R}$ is the utility function on Z in state s .

Thus, (15) can be interpreted as the subjective expected value (over states S with subjective probability π) of the objective expected utility (over outcomes Z given objective probabilities $f(s)$ in state s). We call equation (15) a *state-dependent Anscombe–Aumann representation*. We thus have, as a corollary to Theorem 5 and this discussion, the following representation theorem.

THEOREM 7. *Assume that \succsim satisfies the linearity and Archimedean axioms. Then \succsim has a state-dependent Anscombe–Aumann representation.*

The linearity axiom on \succsim thus encompasses two independence conditions:

1. the *sure-thing principle* as applied to *subjective* uncertainty across different states (that is, linearity implies joint independence over factors, as shown in Section 9.4);
2. the *independence axiom* as applied to *objective* uncertainty within each state (that is, linearity of \succsim implies linearity of the within-state preferences).

The normative arguments that justify these two axioms or principles, which we have already discussed extensively, also justify the linearity axiom in this Anscombe–Aumann framework.

15.3. State independence

The probability measure π is still not uniquely identified because we have state-dependent utility. However, recall from Section 10.5 that the additional assumption of ordinal state independence is enough to obtain state-independent linear utility and thus to pin down the subjective beliefs. The overall representation becomes

$$U(f) = \sum_{s \in S} \pi(s) \sum_{z \in Z} f(s)(z) \times u(z). \quad (16)$$

We call equation (16) an *Anscombe–Aumann representation*.

We summarize this as follows.

THEOREM 8. *Assume that \succsim satisfies the linearity, Archimedean, and ordinal state independence axioms. Then \succsim has an Anscombe–Aumann representation.*

Proof. This follows from Theorem 5 and the preceding discussion. It is also essentially Anscombe and Aumann (1963, Thm. 1) though their axiomatization is a bit different. \square

15.4. Calibration of beliefs

The simplicity of Theorem 8 and the fact that it is a mere application of linear utility masks the way in which beliefs and utilities are disentangled. We illustrate how such calibration happens using ideas that lurk in the proof of the theorem.

For example, consider a less reduced-form version of the story in Section 13.4. Anna chooses between two actions:

- leave* — leaving her current employment to undertake an MBA;
- stay* — staying put.

The relevant outcomes are the three enumerated in Section 13.4: (*e*) no MBA and staying in her current employment; (*m*) bearing the cost of an MBA without then finding a better job; and (*M*) bearing the cost of an MBA and then finding a better job.

The last element in the decision problem is the event E , the set of states in which Anna obtains the better job if she gets an MBA. We take this to be a state or elementary event in the small-worlds model; hence the set of states is $\{s_1, s_2\}$, where s_1 corresponds to event E and s_2 corresponds to event E^c . Therefore the two acts associated with the actions *leave* and *stay* are

$$\begin{aligned} \textit{leave}(s_1) &= M, & \textit{leave}(s_2) &= m; \\ \textit{stay}(s_1) &= e, & \textit{stay}(s_2) &= e. \end{aligned}$$

Whether we have *leave* \succsim *stay* or *stay* \succsim *leave* seems to depend on two separate considerations. How good *Anna feels* the chances of obtaining a better job would have to be in order to make it worthwhile to leave her current employment; and how good *in her opinion* the chances of obtaining a better job actually are. What Anna does when she answers the former question is quantify her *personal*

preference for e relative to m and M . What Anna does when she answers the latter question is quantify her *personal judgment* concerning the relative strengths of the factors that favor and oppose certain events.

In order to calibrate these two considerations, we must assume that she can meaningfully compare *any* horse-race/lottery acts, nor merely the acts *leave* and *stay* available to her in this problem. Thus, she must be able to express preferences over hypothetical acts such as the act

$$g(s_1) = m, \quad g(s_2) = M,$$

(in the state s_2 where she would *not* find a good job if she got an MBA, she gets an MBA and finds a good job!) and the act that yields, in both states, a lottery with equal probability of the three outcomes.

We can first quantify the strength of Anna's personal preference for e relative to m and M by considering the constant acts (lotteries that are not state dependent). That is, we abstract from the subjective uncertainty about the states and consider her preferences over objectively generated lotteries. This is the representation and calibration we covered in Section 13. We thereby let $u(m) = 0$ and $u(M) = 1$ and define $u(e)$ to be such that Anna is indifferent between e and the lottery $u(e)M + (1 - u(e))m$.

To quantify Anna's judgment concerning the likelihood of state s_1 , we let $\pi(s_1)$ be the unique probability for which Anna is indifferent between the act *leave* and the act that leads, in every state, to the lottery with probability $\pi(s_1)$ on M and probability $1 - \pi(s_1)$ on m . The idea is that, given state-independent preferences, the state is simply a randomization device from Anna's point of view: if Anna is indifferent between these two acts, it is because the objective probability $\pi(s_1)$ is the same as Anna's subjective likelihood of state s_1 .

16. Conclusion

Throughout this chapter we have emphasized the link between independence axioms in standard consumer theory, in expected utility theory for decision under objective uncertainty, and in expected utility theory for decision under subjective uncertainty. We contend that the independence axioms have considerable normative appeal in decision under uncertainty.

However, experimental and empirical evidence shows that behavior deviates systematically from these theories, implying that (not surprisingly) such normative theories make for only approximate descriptive models. Furthermore, many authors have disagreed with our claim that the independence axioms are normatively compelling.

There is now a vast literature that has developed generalizations, extensions, and alternatives to expected utility. We will not provide a guide to this literature; doing so would be beyond the scope of our chapter whereas later chapters in this handbook treat it extensively. However, as a transition to those chapters and as a further illustration of the content of the independence axioms, we outline some of the experimental violations.

One of the earliest and best known tests of expected utility is the *common*

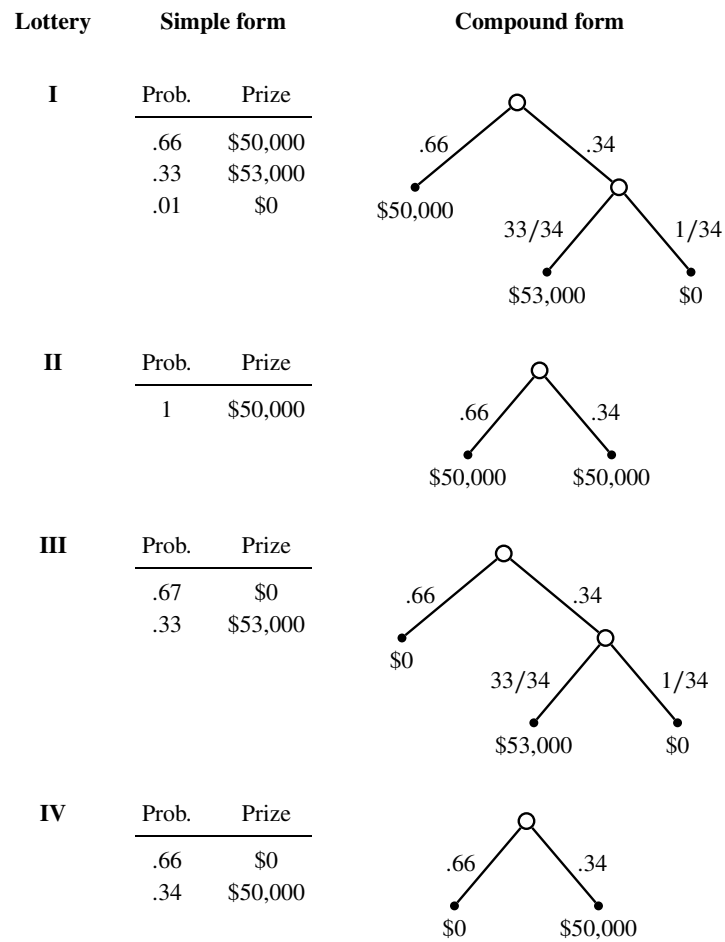


FIGURE 7. Common consequence paradox (Allais paradox): The simple lotteries on the left are the reduced lotteries of the compound lotteries on the right. Preferences $II > I$ and $III > IV$ violate the independence axiom but are common for subjects in decision experiments.




FIGURE 8. Common ratio paradox: The simple lotteries on the left are the reduced lotteries of the compound lotteries on the right. Preferences $II > I$ and $III > IV$ violate the independence axiom but are common for subjects in decision experiments.

consequence paradox, first proposed by Maurice Allais (1953). It is illustrated in Figure 7. Allais conjectured (and found) that most people would prefer lottery II to lottery I but would prefer lottery III to lottery IV (when presented as the simple lotteries on the left). By writing the simple lotteries as the compound lotteries on the right, we can see that such choices violate the independence axiom (Axiom L).

A closely related and frequently observed systematic violation of expected utility theory is the *common ratio paradox* (see Kahneman and Tversky 1979). This is illustrated in Figure 8. Again, the choices of II over I and III over IV are common but violate the independence axiom.

There has been debate about whether these violations are due to bounded rationality or whether the normative model needs adjustment, but there certainly room for better descriptive models than the classic theory reviewed in this chapter (even if, for many applications, the classic theory has proved to be a suitably accurate approximation).

Systematic violations of expected utility theory—observed in choice prob-

lems such as these paradoxes—suggests the following: when altering a lottery by reducing the probability of receiving a given outcome, the portion of the probability we must place on a better outcome (with the remaining portion on a worse outcome) in order to keep the individual indifferent is not *independent* of the lottery with which we began. Yet the independence axiom implies that it is independent. Indeed, for any three outcomes H, M, L , with $H > M > L$, the trade-off for an expected utility maximizer is simply the constant slope of the indifference curves in the simplex of lotteries:

$$\frac{u(M) - u(L)}{u(H) - u(M)}.$$

When assessing the accumulated experimental evidence, Machina (1982) proposed that one could account for these observed systematic violations of expected utility by assuming that this trade-off is *increasing* the “higher up” (in preference terms) in the simplex is the lottery with which one starts. Geometrically, this corresponds to a “fanning out” of the indifference curves in the simplex. Many other patterns have been observed that depend on the size and sign of the payoffs.

In response, several versions of so-called non- or generalized expected utility have axiomatized and analyzed nonlinear representations of preferences over lotteries. These include, among others, rank dependent expected utility of Quiggin (1982) and Yaari (1987), cumulative prospect theory of Tversky and Kahneman (1992) and Wakker and Tversky (1993), betweenness of Dekel (1986) and Chew (1989), and additive bilinear (regret) theories of Loomes and Sugden (1982) and Fishburn (1984).

Another famous experiment, whose results are inconsistent with *subjective* expected utility theory, is the Ellsberg paradox. Daniel Ellsberg (1961) proposed a number of thought experiments to suggest that, in situations with ambiguity about the nature of the underlying stochastic process, preferences over subjectively uncertain acts would not allow for beliefs over the likelihood of events to be represented by a well-defined probability distribution.

One such choice problem involves an urn from which a ball will be drawn. Anna knows there are 90 balls in total, of which 30 are red. However, the only information she has about the remaining 60 balls is that some are black and some are white—she is not told the actual proportions. Consider two choice problems.

1. A choice between (f) an act that pays \$100 dollars if the ball drawn is red and nothing if it black or white, and (g) an act that pays \$100 dollars if the ball drawn is black and nothing if it is red or white.
2. A choice between (f') an act that pays \$100 dollars if the ball drawn is red or white and nothing if it is black, and (g') an act that pays \$100 dollars if the ball drawn is black or white and nothing if it is red.

Ellsberg conjectured that many individuals would be averse to ambiguity in the sense that they would prefer to bet on “known” rather than “unknown” odds. In this example, they would strictly prefer the bet on red to the bet on black in the first problem ($f > g$)—indicating a subjective belief that black is *less* likely than red—but then prefer the bet on black or white to the bet on red or white in the second problem ($g' > f'$)—indicating a subjective belief that black is *more* likely

than red. As indicated, such a preference pattern is inconsistent with beliefs being represented by a well-defined probability measure.

In response, models have been developed in which beliefs are represented by multiple measures and/or nonadditive capacities (which is a generalization of a probability measure). Examples are Gilboa and Schmeidler (1989) and Schmeidler (1989).

We have mentioned only a small sample of critiques of classic expected utility theory reviewed and of the extensions to that theory. This theme is developed further in other chapters of this handbook.

References

- Allais, Maurice (1953). La psychologie de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine. *Econometrica*, 21, 503–546.
- Anscombe, F.J. and R.J. Aumann (1963). A Definition of Subjective Probability. *Annals of Mathematical Statistics*, 34, 199–205.
- Arrow, Kenneth J. (1959). Rational Choice Functions and Orderings. *Economica*, 26, 121–127.
- Bernoulli, Daniel (1738). Specimen Theoriae Novae de Mensura Sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, 5, 175–192.
- Birkhoff, Garrett (1948). *Lattice Theory*. New York: American Mathematical Society.
- Cantor, Georg (1895). Beiträge zur Begründung der transnieten Mengenlehre I. *Mathematische Annalen*, 46, 481–512.
- Chew, Soo Hong (1989). Axiomatic Utility Theories with the Betweenness Property. *Annals of Operations Research*, 19, 273–298.
- de Finetti, Bruno (1931). Sul Significato Soggettivo della Probabilità. *Fundamenta Mathematicae*, 17, 298–329.
- Debreu, Gerard (1954). Representation of a Preference Ordering by a Numerical Function. In R. M. Thrall, C. H. Coombs, and R. L. Davis (Eds.), *Decision Processes*, pp. 159–165. New York: Wiley.
- Debreu, Gerard (1960). Topological Methods in Cardinal Utility Theory. In K. J. Arrow, S. Karlin, and P. Suppes (Eds.), *Mathematical Methods in the Social Sciences, 1959*, pp. 16–26. Stanford University Press.
- Debreu, Gerard (1964). Continuity Properties of Paretian Utility. *International Economic Review*, 5, 285–293.
- Dekel, Eddie (1986). An Axiomatic Characterization of Preferences under Uncertainty: Weakening the Independence Axiom. *Journal of Economic Theory*, 40, 304–318.
- Ellsberg, Daniel (1961). Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 75, 643–669.
- Fishburn, Peter (1970). *Utility Theory for Decision Making*. New York: John Wiley & Sons.
- Fishburn, Peter (1984). Dominance in SSB Utility Theory. *Journal of Economic Theory*, 34, 130–148.
- Fishburn, Peter and Peter Wakker (1995). The Invention of the Independence Condition for Preferences. *Management Science*, 41, 1130–1144.
- Gilboa, Itzhak and David Schmeidler (1989). Maxmin Expected Utility with a Non-Unique Prior. *Journal of Mathematical Economics*, 18, 141–153.
- Hammond, Peter J. (1988). Consequentialist Foundations for Expected Utility. *Theory and Decision*, 25, 25–78.

- Herstein, I.N. and John Milnor (1953). An Axiomatic Approach to Measurable Utility. *Econometrica*, 21, 291–297.
- Jensen, Niels-Erik (1967). An Introduction to Bernoullian Utility Theory, I, II. *Swedish Journal of Economics*, 69, 163–183, 229–247.
- Kahneman, Daniel and Amos Tversky (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47, 263–291.
- Karni, Edi (1985). *Decision Making under Uncertainty: The Case of State-Dependent Preferences*. Cambridge, MA: Harvard University Press.
- Krantz, David H., R. Duncan Luce, Patrick Suppes, and Amos Tversky (1971). *Foundations of Measurement*, volume I. Additive and Polynomial Representations. New York: Academic Press.
- Loomes, Graham and Robert Sugden (1982). Regret Theory: An Alternative Theory of Rational Choice under Uncertainty. *Economic Journal*, 92, 805–824.
- Machina, Mark J. (1982). ‘Expected Utility’ Analysis without the Independence Axiom. *Econometrica*, 50, 277–323.
- Quiggin, John (1982). A Theory of Anticipated Utility. *Journal of Economic Behavior and Organization*, 3, 323–343.
- Ramsey, Frank P. (1931). Truth and Probability. In *Foundations of Mathematics and other Logical Essays*. London: K. Paul, Trench, Trubner, and Co.
- Samuelson, Paul A. (1938). A Note on the Pure Theory of Consumer’s Behaviour. *Economica*, 5, 61–71.
- Savage, Leonard J. (1954). *The Foundations of Statistics*. New York: Wiley.
- Schmeidler, David (1989). Subjective Probability and Expected Utility without Additivity. *Econometrica*, 57, 571–587.
- Strotz, Robert H. (1959). The Utility Tree—A Correction and Further Appraisal. *Econometrica*, 27, 482–488.
- Tversky, Amos and Daniel Kahneman (1992). Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, 5, 297–323.
- von Neumann, John and Oskar Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Wakker, Peter (1988). The Algebraic versus the Topological Approach to Additive Representations. *Journal of Mathematical Psychology*, 32, 421–435.
- Wakker, Peter (1989). *Additive Representations of Preferences: A New Foundation of Decision Analysis*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Wakker, Peter and Amos Tversky (1993). An Axiomatization of Cumulative Prospect Theory. *Journal of Risk and Uncertainty*, 7, 147–176.
- Wakker, Peter and Horst Zank (1999). State Dependent Expected Utility for Savage’s State Space. *Mathematics of Operations Research*, 24, 8–34.
- Yaari, Menahem E. (1987). The Dual Theory of Choice under Risk. *Econometrica*, 55, 95–15.