CHAPTER 8

Choice Under Uncertainty

In this chapter we turn to models of choice under uncertainty. We consider an agent who makes choices without fully knowing the consequences of those choices, and focus on models in which uncertainty can be quantified and formulated probabilistically. The most important such model is, of course, expected utility.

8.1 OBJECTIVE PROBABILITY

There are times when probabilities can be thought to be objective and known, or observable. This is the case, for example, when outcomes are randomized according to some known physical device—such as a game in a casino, or a randomization device used by an experimenter in the laboratory.

We consider two basic environments. In one the primitive objects of choice are lotteries. In the other, the objects of choice are state-contingent consumption.

8.1.1 Notation

Let X be a finite set. We denote by $\Delta(X) = \{p \in \mathbf{R}^X : p \geq 0; \sum_{x \in X} p(x) = 1\}$ the set of all probability distributions over X.

8.1.2 Choice over lotteries

Given is a finite set X of possible *prizes*. $\Delta(X)$ is the set of all *lotteries* over X. We imagine an agent who chooses a lottery. The agent understands that there is uncertainty over the realization of the lottery: over which prize the lottery will result in. But the probabilities specified in the lottery are accurate (or at least useful) representations of that uncertainty.

We investigate a very basic result on revealed preference in this environment.

An *expected utility preference* $\succeq$ is a binary relation for which there exists $u : X \to \mathbf{R}$ such that for all $p,q \in \Delta(X)$,

$$p \succeq q \text{ iff } \sum_{x\in X} p(x)u(x) \geq \sum_{x\in X} q(x)u(x).$$

The classical axiomatization of expected utility preferences relies on the *independence axiom* of decision theory; namely, that for all $p,q,r \in \Delta(X)$ and all $\alpha \in (0,1]$, $p \succeq q$ iff $\alpha p + (1-\alpha)r \succeq \alpha q + (1-\alpha)r$.

Most experimental studies refuting the expected utility model are direct refutations of the independence axiom. The best-known such refutation is through a thought experiment, known as the Allais paradox. Instead of setting up a thought experiment, we are going to assume that we are given data on choices among pairs of lotteries.

The data can be organized into a revealed preference pair $\langle \succeq^R, \succ^R \rangle$, where each of $\succeq^R$ and $\succ^R$ are finite sets. The idea is that an agent makes "binary choices:" choices from budgets with two alternatives, and these choices define $\langle \succeq^R, \succ^R \rangle$. We ask when there exists an expected utility preference $\succeq$ such that for all $p,q \in \Delta(X)$, $p \succeq^R q$ implies $p \succeq q$ and $p \succ^R q$ implies $p \succ q$.

The following example demonstrates that observed data can be incompatible with the expected utility model without directly violating the independence axiom.

Example 8.1 *Let $X = \{x,y,z\}$, and consider the rankings:* $(1,0,0) \succ^R (1/3,1/3,1/3)$*,* $(0,1,0) \succ^R (1/3,1/3,1/3)$*,* $(0,0,1) \succ^R (1/3,1/3,1/3)$*. This is clearly incompatible with the expected utility model: namely, the rankings would imply that $u(x),u(y),u(z) > \frac{u(x)+u(y)+u(z)}{3}$, which is impossible. However, there is no direct refutation of the independence axiom. There is, instead, a refutation of the joint hypotheses of the independence axiom and transitivity. (Note, incidentally, that the example is not a direct refutation of transitivity either.) Our aim is to uncover all implications of the joint hypotheses of independence and transitivity for finite datasets.*

Theorem 8.2 *Suppose that each of $\succeq^R$ and $\succ^R$ are finite, and without loss of generality let $\succeq^R = \{(p_i,q_i)\}_{i=1}^K$ and $\succ^R = \{(p_i,q_i)\}_{i=K+1}^L$. There is an expected utility preference $\succeq$ such that for all $p,q \in \Delta(X)$, $p \succeq^R q$ implies $p \succeq q$ and $p \succ^R q$ implies $p \succ q$ iff there is no $\lambda \in \Delta(L)$ for which $\lambda(\{K+1,\ldots,L\}) > 0$ and $\sum_{i=1}^L \lambda_i p_i = \sum_{i=1}^L \lambda_i q_i$.*

It is worth remarking on what Theorem 8.2 says. The $\lambda \in \Delta(L)$ can be understood as a lottery over lotteries (a "first stage" lottery), which is compounded either with the lotteries $\{p_i\}_{i=1}^L$ or the lotteries $\{q_i\}_{i=1}^L$. When λ is compounded with the lotteries $\{p_i\}_{i=1}^L$, the reduced lottery is $\sum_{i=1}^L \lambda_i p_i$, and when compounded with $\{q_i\}_{i=1}^L$, it is $\sum_{i=1}^L \lambda_i q_i$. The condition in Theorem 8.2 requires that, for weights λ, it is impossible that the corresponding p lotteries are preferred *ex-post*, yet *ex-ante* the two compound lotteries are identical. Figure 8.1 illustrates the condition in Theorem 8.2.

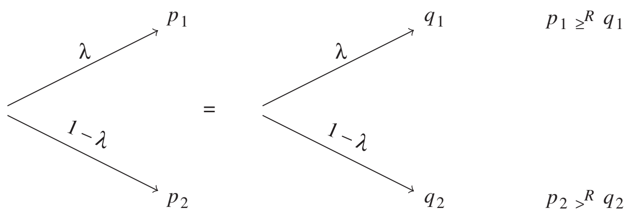

Fig. 8.1 An illustration of Theorem 8.2.

Proof. The proof is almost a direct translation of Lemma 1.12 to this environment. Namely, we seek the existence of $u \in \mathbf{R}^X$ such that for all $i = 1,\ldots,K$, $(p_i - q_i) \cdot u \geq 0$, and for all $i = K+1,\ldots,L$, $(p_i - q_i) \cdot u > 0$. Non-existence of u is therefore equivalent to existence of $\lambda \in \mathbf{R}^L_+$ such that for some $i \in \{K+1,\ldots,L\}$, $\lambda_i > 0$ and $\sum_{i=1}^L (p_i - q_i) = 0$. Since $\sum_{i=1}^L \lambda_i > 0$, we can normalize λ so that $\sum_{i=1}^L \lambda_i = 1$.

8.1.3 State-contingent consumption

Many applications of choice under uncertainty in economic models involve a state-contingent environment. Suppose that there is a finite set Ω of *states of the world*. A state-contingent consumption bundle is modeled as a vector in $\mathbf{R}^\Omega_+$. An agent chooses $x \in \mathbf{R}^\Omega_+$: If the state of the world is $\omega \in \Omega$ then the agent obtains a monetary payment of x_ω. The vectors $x \in \mathbf{R}^\Omega_+$ are referred to as *monetary acts* in decision theory.

The focus of our discussion will be expected utility theory with risk aversion. Consider an agent with a known *prior probability measure* $\pi \in \Delta(\Omega)$ describing her beliefs over the possible states of the world. We suppose that for all $\omega \in \Omega$, $\pi_\omega > 0$. The prior is known. This means that it is observable; possibly it has been induced by some experimental design (for example, experiments in economics often use a randomization device).

Expected utility theory says that the choice of $x \in \mathbf{R}^\Omega_+$ is determined by maximizing a utility function of the form

$$U(x) = \sum_{\omega \in \Omega} \pi_\omega u(x_\omega),$$

where $u : \mathbf{R}_+ \to \mathbf{R}$ is a strictly monotonically increasing and concave function. The function u is a utility function over money, and the concavity of u means that the agent in question is risk averse.

Suppose we are given data on the behavior of our agent in the market. When faced with prices $p = (p_\omega)_{\omega \in \Omega} \in \mathbf{R}^\Omega_{++}$ and an income $I > 0$, the agent chooses x to maximize $U(x)$ over the x that she can afford. It is important to emphasize what the meaning of prices are here. In the development of the theory, we take prices for state-contingent consumption as given, but to have access to such prices in real data requires the relevant financial markets to be complete.[1]

A dataset is a collection $(x^k, p^k)_{k\in K}$ of pairs of a consumption $x^k \in \mathbf{R}^\Omega_+$ chosen at prices $p^k \in \mathbf{R}^\Omega_{++}$ and budget $p^k \cdot x^k$. Here K denotes the set $\{1,\ldots,K\}$, an instance of inconsequential notational abuse.

We are given a probability distribution $\pi \in \Delta(\Omega)$. A function $u : \mathbf{R}_+ \to \mathbf{R}$ *weakly expected-utility rationalizes* a dataset $(x^k, p^k)_{k\in K}$ if u is strictly increasing and concave, and if $p^k \cdot y \le p^k \cdot x^k$ implies that

$$\sum_{\omega\in\Omega} \pi_\omega u(y_\omega) \le \sum_{\omega\in\Omega} \pi_\omega u(x^k_\omega)$$

for all $k \in K$. We should emphasize that:

I) The prior π is given and known. We are interested in testing whether the agent behaves according to expected utility with respect to prior π, rather than any other (possibly subjective) prior.

II) The exercise is restricted to concave utility functions. Concavity means that the agent is risk averse.

The existence of a known prior π allows us to compute *risk-neutral prices*, defined as follows: for $k \in K$ and $\omega \in \Omega$, let

$$\rho^k_\omega = \frac{p^k_\omega}{\pi_\omega}.$$

Risk-neutral prices turn out to be the relevant prices one needs to use to test for expected utility theory.

We can gain some intuition for how the problem can be solved by considering the case when $u : \mathbf{R}_+ \to \mathbf{R}$ is differentiable. The first-order condition for utility maximization demands that, for any k and ω, we have $\pi_\omega u'(x^k_\omega) = \lambda^k p^k_\omega$, where λ^k is the Lagrange multiplier for the maximization problem in which x^k is chosen, and we have assumed an interior optimum. Using the definition of risk-neutral prices, we obtain that $u'(x^k_\omega) = \lambda^k \rho^k_\omega$. It follows from concavity of u, then, that $x^k_\omega > x^k_{\omega'}$ implies $\rho^k_\omega/\rho^k_{\omega'} \le 1$. This property can be thought of as *downward-sloping demand*.

More generally, if we have $x^k_\omega > x^{k'}_{\omega'}$ and $x^{k'}_{\omega''} > x^k_{\omega'''}$ then we obtain that

$$1 \ge \frac{u'(x^k_\omega)}{u'(x^{k'}_{\omega'})} \cdot \frac{u'(x^{k'}_{\omega''})}{u'(x^k_{\omega'''})} = \frac{\rho^k_\omega}{\rho^{k'}_{\omega'}} \cdot \frac{\rho^{k'}_{\omega''}}{\rho^k_{\omega'''}},$$

as u is concave, and Lagrange multipliers cancel out. The implication is not that two pairs of quantities and prices are inversely related, but in some sense that $(x^k_\omega, x^{k'}_{\omega'})$ and $(x^{k'}_{\omega''}, x^k_{\omega'''})$ are "inversely related" to $(\rho^k_\omega, \rho^{k'}_{\omega'})$ and $(\rho^{k'}_{\omega''}, \rho^k_{\omega'''})$.

The preceding calculation suggests that one should consider sequences of pairs $(x^{k_i}_{\omega_i}, x^{k'_i}_{\omega'_i})^n_{i=1}$ for which each k appears in k_i (on the left of the pair) the

[1] Prices p can be calculated (uniquely) from an array of asset prices under two assumptions: that the markets are complete and that there is no arbitrage.

same number of times it appears in k'_i (on the right). Such sequences will be called *balanced*.

A dataset $\{(x^k, p^k)\}_{k=1}^K$ satisfies the *Strong Axiom of Revealed Objective Expected Utility (SAREU)* if, for any balanced sequence of pairs $(x^{k_i}_{\omega_i}, x^{k'_i}_{\omega'_i})_{i=1}^n$ with the property that $x^{k_i}_{\omega_i} > x^{k'_i}_{\omega'_i}$ for all i, the product of relative risk-neutral prices satisfies

$$\prod_{i=1}^n \frac{\rho^{k_i}_{\omega_i}}{\rho^{k'_i}_{\omega'_i}} \leq 1.$$

It is possible to write SAREU so that it rules out certain kinds of cycles. We use the syntax above to make the comparison with subjective probabilities in Section 8.2.2 (below) easier.

Theorem 8.3 *The following statements are equivalent:*

I) $\{(x^k, p^k)\}_{k=1}^K$ *is weakly expected-utility rationalizable.*
II) *For all $k \in K$ and $s \in S$ there exist $\lambda^k > 0$ and $v^k_\omega > 0$ such that*

$$\pi_\omega v^k_\omega = \lambda^k p^k_\omega$$

and $x^k_\omega > x^{k'}_{\omega'}$ implies that $v^k_\omega \leq v^{k'}_{\omega'}$.
III) $\{(x^k, p^k)\}_{k=1}^K$ *satisfy SAREU.*

Remark 8.4 Statement (II) asserts the existence of Afriat inequalities for this problem. The numbers v^k_ω are meant to be the marginal utilities (for money) at the quantity x^k_ω, and λ^k is meant to be the Lagrange multiplier.

We could equivalently have written Statement (II) as: For all $k \in K$ and $\omega \in \Omega$ there exist $\lambda^k > 0$ and $u^{k,\omega}$ such that

$$u^{k,\omega'} \leq u^{l,\omega} + \lambda^l \frac{p^l_\omega}{\pi_\omega}[x^k_{\omega'} - x^l_\omega],$$

for all $k, l \in K$ and all $\omega, \omega' \in \Omega$.

Proof. To begin with, we establish the equivalence between statements (II) and (III), the more interesting aspect of the proof of Theorem 8.3. The reader should note the similarities between the construction in the proof that (III) implies (II), and the construction in Theorem 1.9. The proof of the equivalence of (I) and (II) is standard, and is included below for completeness' sake.

Let $Y = \{x^k_\omega : k \in K; \omega \in \Omega\}$; enumerate the elements of Y in increasing order, as follows:

$$y_1 < y_2 < \ldots < y_n.$$

First we show that (III) implies (II). Suppose that $\{(x^k, p^k)\}_{k=1}^K$ satisfies SAREU. We shall construct a solution to the system

$$\pi_\omega v^k_\omega = \lambda^k p^k_\omega \tag{8.1}$$

$$x^k_\omega > x^{k'}_{\omega'} \Longrightarrow v^k_\omega \leq v^{k'}_{\omega'}, \tag{8.2}$$

where $\lambda^k > 0$ and $v^k_\omega > 0$ for each k, thereby proving (II).

Let k_0 be such that $x^{k_0}_\omega = y_n$ for some ω, and such that

$$\rho^{k_0}_\omega = \max\{\rho^{k'}_{\omega'} : x^{k'}_{\omega'} = y_n\}.$$

For $k, l \in K$, define $\eta(k,l)$ as follows. Let

$$\eta(k,l) = \max\left\{ \frac{\rho^k_\omega}{\rho^l_{\omega'}} : x^k_\omega > x^l_{\omega'} \right\}$$

if there is $\omega, \omega' \in \Omega$ such that $x^k_\omega > x^l_{\omega'}$, and $\eta(k,l) = 0$ otherwise.

Note that SAREU implies that

$$\eta(k_1,k_2)\eta(k_2,k_3)\cdots\eta(k_m,k_1) \leq 1.$$

Therefore, for any sequence $k_1,\ldots,k_m \in K$, in which $k_i = k_j$ for $i < j$, we have that

$$\begin{aligned} &\eta(k_0,k_1)\eta(k_1,k_2)\cdots\eta(k_{i-1},k_i)\eta(k_i,k_{i+1})\cdots\eta(k_{j-1},k_j)\eta(k_j,k_{j+1})\cdots \\ &\eta(k_{m-1},k_m) \leq \eta(k_0,k_1)\eta(k_1,k_2)\cdots\eta(k_{i-1},k_i)\eta(k_j,k_{j+1})\cdots\eta(k_{m-1},k_m). \end{aligned} \tag{8.3}$$

Let $\lambda^l = 1$ whenever l is such that there is $\omega \in \Omega$ such that $x^l_\omega = y_n$. For any other l, let

$$\lambda^l = \max\{\eta(k_0,k_1)\eta(k_1,k_2)\cdots\eta(k_m,l) : k_1,\ldots,k_m \in K \text{ and } m \geq 0.\}$$

Observe that λ^l is well defined, as Equation (8.3) implies that one can restrict attention to a finite set of sequences $k_1,\ldots,k_m$; observe also that $\lambda^l > 0$.

The definition of $(\lambda^l)_{l\in K}$ implies that, when $y_n > x^k_\omega > x^l_{\omega'}$ we have

$$\lambda^l \geq \eta(k,l)\lambda^k \geq \frac{\rho^k_\omega}{\rho^l_{\omega'}}\lambda^k; \tag{8.4}$$

and when $y_n = x^k_\omega > x^l_{\omega'}$ we have that, for some $\hat{\omega}$,

$$\lambda^l \geq \eta(k_0,l) \geq \frac{\rho^{k_0}_{\hat{\omega}}}{\rho^l_{\omega'}} \geq \frac{\rho^k_\omega}{\rho^l_{\omega'}}\lambda^k;$$

where the first inequality is by definition of λ^l; the second by definition of η (as $y_n = x^{k_0}_{\hat{\omega}} > x^l_{\omega'}$); the third by definition of k_0, and where we have used that $\lambda^k = 1$ in this case. Therefore we have that $\lambda^l \geq \frac{\rho^k_\omega}{\rho^l_{\omega'}}\lambda^k$ holds whenever $x^k_\omega > x^l_{\omega'}$, regardless of whether $x^k_\omega = y_n$ or $x^k_\omega < y_n$.

If we let $v^l_\omega = \lambda^l \rho^l_\omega$ for all $l \in K$ and $\omega \in \Omega$, then we have solutions to system (8.1)–(8.2).

Conversely, suppose that (II) is true. Let there be strictly positive numbers v^k_ω, λ^k, for $\omega \in \Omega$ and $k = 1,\ldots,K$, solving system (8.1)–(8.2).

Let $(x^{k_i}_{\omega_i}, x^{k'_i}_{\omega'_i})^n_{i=1}$ be a sequence of pairs under the assumptions of SAREU. Since $x^{k_i}_{\omega_i} > x^{k'_i}_{\omega'_i}$ we have that $v^{k_i}_{\omega_i} \leq v^{k'_i}_{\omega'_i}$. Hence

$$1 \leq \prod_{i=1}^{n} \frac{v^{k'_i}_{\omega'_i}}{v^{k_i}_{\omega_i}} = \prod_{i=1}^{n} \frac{\lambda^{k'_i} \rho^{k'_i}_{\omega'_i}}{\lambda^{k_i} \rho^{k_i}_{\omega_i}} = \prod_{i=1}^{n} \frac{\rho^{k'_i}_{\omega'_i}}{\rho^{k_i}_{\omega_i}},$$

as each k appears as k_i the same number of times it appears as k'_i in the sequence, and therefore the λ^{k_i} and $\lambda^{k'_i}$ cancel out. So the data satisfy SAREU, and thus we establish (III).

We shall now prove that (I) implies (II). Let $(x^k, p^k)^K_{k=1}$ be weakly rationalizable by an expected utility preference with prior π. Let $u : \mathbf{R}_+ \to \mathbf{R}$ be a strictly increasing and concave rationalizing utility function. We can then consider the first-order conditions for a maximizing utility: see, for example, Theorem 28.3 of Rockafellar (1997) for a formulation that does not require u to be smooth. The first-order conditions say that there are numbers $\lambda^k \geq 0$, $k = 1, \ldots, K$ such that if we let

$$v^k_\omega = \frac{\lambda^k p^k_\omega}{\pi_\omega}$$

then $v^k_\omega \in \partial u(x^k_\omega)$ if $x^k_\omega > 0$, and there is $\underline{w} \in \partial u(x^k_\omega)$ with $v^k_\omega \geq \underline{w}$ if $x^k_\omega = 0$. In fact, since u is strictly increasing it is easy to see that $\lambda^k > 0$, and therefore $v^k_\omega > 0$.

By the concavity of u, and the consequent monotonicity of $\partial u(x^k_\omega)$ (Theorem 1.9), if $x^k_\omega > x^{k'}_{\omega'} > 0$, $v^k_\omega \in \partial u(x^k_\omega)$, and $v^{k'}_{\omega'} \in \partial u(x^{k'}_{\omega'})$, then $v^k_\omega \leq v^{k'}_{\omega'}$. If $x^k_\omega > x^{k'}_{\omega'} = 0$, then $\underline{w} \in \partial u(x^{k'}_{\omega'})$ with $v^{k'}_{\omega'} \geq \underline{w}$. So $v^k_\omega \leq \underline{w} \leq v^{k'}_{\omega'}$.

Next, we show that (II) implies (I). Suppose that the numbers v^k_ω, λ^k, for $s = 1, \ldots, S$ and $k = 1, \ldots, K$, are as in (II).

Let

$$\underline{y}_i = \min\{v^k_\omega : x^k_\omega = y_i\} \text{ and } \bar{y}_i = \max\{v^k_\omega : x^k_\omega = y_i\}.$$

Let $z_i = (y_i + y_{i+1})/2$, $i = 1, \ldots, n-1$; $z_0 = 0$, and $z_n = y_n + 1$. Let f be a correspondence defined as follows:

$$f(z) = \begin{cases} [\underline{y}_i, \bar{y}_i] & \text{if } z = y_i, \\ \max\{\bar{y}_i : z < y_i\} & \text{if } y_n > z \text{ and } \forall i (z \neq y_i), \\ \underline{y}_n / 2 & \text{if } y_n < z. \end{cases}$$

By the assumptions placed on v^k_ω, and by construction of f, $y < y'$, $v \in f(y)$ and $v' \in f(y')$ imply that $v' \leq v$. Then the correspondence f is monotone, and there exists a concave function u for which $f(z) \subseteq \partial u(z)$ (see Corollary 1.10 of Theorem 1.9). Given that $v^k_\omega > 0$ for all k and ω, all the elements in the range of f are positive, and therefore u is a strictly increasing function.

Finally, for all (k,s), $p^k_\omega/\pi_\omega = v^k_\omega \in \partial u(v^k_\omega)$ and therefore the first-order conditions to a maximum choice of x hold at x^k_ω. Since u is concave the first-order conditions are sufficient. The data are therefore rationalizable.

8.2 SUBJECTIVE PROBABILITY

The expected utility model of Section 8.1 assumes that an agent's behavior can be represented probabilistically, and in Section 8.1 probabilities are in fact observable. However, in most situations of interest to economists, probabilities are not observable. Economic agents are assumed that assign subjective probabilities to the states of the world. It is important to understand when agents' behavior is consistent with the use of subjective probabilities.

8.2.1 The Epstein Test

One of the most basic questions in the theory of choice under uncertainty is whether individuals perceive uncertainty probabilistically. One way of formalizing this idea is due to Machina and Schmeidler (1992), and is called probabilistic sophistication. A preference $\succeq$ over $\mathbf{R}^\Omega_+$ is said to be *probabilistically sophisticated* if there is a probability measure π on Ω such that for all $x,y \in \mathbf{R}^\Omega$, if the random variable x first-order stochastically dominates y on the probability space (Ω,π), then $x \succeq y$; and $x \succ y$ when the first-order stochastic dominance is strict.[2]

Consider a dataset $\{(x^k,p^k)\}^K_{k=1}$, as in 8.1. The dataset is *weakly rationalizable by a probabilistically sophisticated preference* if there is a probabilistically sophisticated preference $\succeq$ such that $x^k \succeq y$, for all $y \in \mathbf{R}^\Omega_+$ with $p^k \cdot y \leq p^k \cdot x^k$.

The following test of the probabilistic sophistication hypothesis is due to Larry Epstein. The idea is as follows. For two states, ω and ω', what type of behavior could reveal that $\pi_\omega \geq \pi_{\omega'}$? Epstein's idea was that if prices in state ω are higher than they are in state ω', and the individual demands more in state ω, then if the probabilistic sophistication hypothesis were true, the only reason this could occur would be if $\pi_\omega > \pi_{\omega'}$.

Theorem 8.5 *Suppose the dataset $\{(x^k,p^k)\}^K_{k=1}$ contains observations l,m for which $p^l_\omega > p^l_{\omega'}$ and $p^m_\omega \leq p^m_{\omega'}$, and $x^l_\omega > x^l_{\omega'}$ and $x^m_{\omega'} > x^m_\omega$. Then there is no probabilistically sophisticated preference which weakly rationalizes the data.*

Proof. Suppose by way of contradiction that there is a probabilistically sophisticated rationalization with associated probability π. We know that either $\pi_\omega > \pi_{\omega'}$ or $\pi_{\omega'} \geq \pi_\omega$. In the first case, because $p^m_\omega \leq p^m_{\omega'}$, we know that

$$p^m_\omega x^m_{\omega'} + p^m_{\omega'} x^m_\omega \leq p^m_\omega x^m_\omega + p^m_{\omega'} x^m_{\omega'}.$$

[2] We say a random variable x first-order stochastically dominates y on probability space (Ω,π) if for all $a \in \mathbf{R}$, $\pi(\{\omega \in \Omega : X(\omega) \geq a\}) \geq \pi(\{\omega \in \Omega : Y(\omega) \geq a\})$.

But the bundle which results by switching consumption in states ω and ω' in x^m strictly first-order stochastically dominates x^m, a contradiction. So it must be the case that $\pi_{\omega'} \geq \pi_\omega$. Now, we have

$$p^l_\omega x^l_{\omega'} + p^l_{\omega'} x^l_\omega < p^l_\omega x^l_\omega + p^l_{\omega'} x^l_{\omega'},$$

so the bundle which results from switching consumption in states ω and ω' for bundle x^l first-order stochastically dominates x^l, and is strictly cheaper. By increasing consumption in every state, there is a bundle which strictly first-order stochastically dominates x^l and is feasible at prices p^l, a contradiction to the fact that x^l is demanded.

8.2.2 Subjective expected utility

The benchmark model of decisions under uncertainty is the model of *subjective expected utility* (SEU). This model postulates that agents' choices are governed by expected utility calculations, as in Section 8.1.3, but where the prior π is not given, or observable. Instead, the agents' choices are *as if* there were some prior, and some utility function, that could explain them.

As in Section 8.1.3, a dataset is a collection $\{(x^k, p^k)_{k \in K}\}$ of pairs of a consumption $x^k \in \mathbf{R}^\Omega_+$ chosen at prices $p^k \in \mathbf{R}^\Omega_{++}$ and income $p^k \cdot x^k$.

A utility function $u : \mathbf{R}_+ \to \mathbf{R}$, together with a prior $\pi \in \Delta(\Omega)$, *weakly subjective expected-utility rationalizes* a dataset $(x^k, p^k)_{k \in K}$ if u is strictly increasing and concave, and if $p^k \cdot y \leq p^k \cdot x^k$ implies that

$$\sum_{\omega \in \Omega} \pi_\omega u(y_\omega) \leq \sum_{\omega \in \Omega} \pi_\omega u(x^k_\omega),$$

for all $k \in K$. We say that such a dataset is weakly SEU rationalizable.

Viewed in this light, it is clear that the SEU model is a special case of both the additive separable model considered in Section 4.2.2, and the probabilistically sophisticated model of Section 8.2.1.

Now one can reason along the same lines of 8.1.3 to obtain an implication from quantities on prices. The complication is that, since the prior π is not observable, one cannot subsume probabilities into risk-neutral prices. Instead, unknown probabilities need to be accounted for in the analysis. The details are omitted; but suffice it to say that we need the same notion of balancedness as in 8.1.3. In fact, we need more: say that a sequence of pairs $(x^{k_i}_{\omega_i}, x^{k'_i}_{\omega'_i})^n_{i=1}$ is *doubly balanced* if it is balanced and if, moreover, each $\omega \in \Omega$ appears as ω_i (on the left of the pair) the same number of times it appears as ω'_i (on the right).

A dataset $\{(x^k, p^k)\}^K_{k=1}$ satisfies the *Strong Axiom of Revealed Subjective Expected Utility (SARSEU)* if, for any sequence of pairs $(x^{k_i}_{\omega_i}, x^{k'_i}_{\omega'_i})^n_{i=1}$ for which

$x_{\omega_i}^{k_i} > x_{\omega_i'}^{k_i'}$ for all i, the product of relative prices satisfies

$$\prod_{i=1}^{n} \frac{p_{\omega_i}^{k_i}}{p_{\omega_i'}^{k_i'}} \leq 1.$$

SAREU in 8.1.3 can be written as ruling out certain kinds of cycles. With SARSEU this is not possible because the axiom involves, in a sense, pairs of cycles (one cycle in k's and one in ω's).

Theorem 8.6 *The following statements are equivalent:*

I) $\{(x^k, p^k)\}_{k=1}^K$ *is weakly SEU rationalizable.*
II) *For all* $k = \{1, \ldots, K\}$*, there exist* $\lambda^k > 0$ *and* $\pi \in \Delta(\Omega)$ *such that for all* $\omega \in \Omega$*,* $\pi_\omega > 0$*, and for each pair* k, ω*, there exists* $u^{k,\omega}$ *such that for all* k, l *and all* ω, ω'*,*

$$u^{k,\omega'} \leq u^{l,\omega} + \lambda^l \frac{p_\omega^l}{\pi_\omega}[x_{\omega'}^k - x_\omega^l].$$

III) $\{(x^k, p^k)\}_{k=1}^K$ *satisfies SARSEU.*

We will not offer a proof of Theorem 8.6, but note that it rests on familiar ideas. Basically, (II) in Theorem 8.6 are the relevant Afriat inequalities for the problem at hand. We only care to construct a utility index for a single commodity, which should be the same across states after rescaling; π_ω here acts as a scaling factor. Importantly, these Afriat inequalities define a nonlinear polynomial system in the variables u, λ, π, which presents a significant complication. The proof that SARSEU characterizes SEU rationalizability rests on linearizing the Afriat inequalities, and then using an approximation argument.

8.3 COMPLETE CLASS RESULTS

The above discussion has focused on consumption data, but one could ask the same type of question in an abstract environment of choice. We now assume as given a single observation of a choice in an abstract environment: the set X is some abstract payoff space; Ω is a finite set of states; and the objects of choice are *acts*, mappings $f : \Omega \to X$. The set of all acts is X^Ω.

Choice is modeled by a preference relation $\succeq$ over acts. The important aspect of $\succeq$ will be that it has a maximal element in some "budget," or set of available acts, $\mathcal{F}$; so we know that there is $f^* \in \mathcal{F}$ such that $f^* \succ g$ for all $g \in \mathcal{F} \setminus \{f^*\}$. The existence of f^* can be obtained from a single observation of a choice at $\mathcal{F}$.[3]

[3] Of course, information on the other comparisons contained in $\succeq$ would have to entail additional observations.

A *constant act* is one whose payoff is independent of the state. We identify an element $x \in X$ with the constant act that takes the value x for all states.

A *null state* is a state ω with the property that the agent does not care about what obtains on ω. Formally, $\omega \in \Omega$ is *null* if for all $x, y \in X$ and $f \in X^{\Omega}$, $x\omega f \sim y\omega f$, where $x\omega f$ denotes the act which pays x if ω obtains and f otherwise.

A preference $\succeq$ on X^{Ω} satisfies *monotonicity* if for all $f, g \in X^{\Omega}$, $f(\omega) \succeq g(\omega)$ for all $\omega \in \Omega$ implies $f \succeq g$, with strict preference if there is non-null $\omega \in \Omega$ for which $f(\omega) \succ g(\omega)$.

In this context, a preference $\succeq$ is a *subjective expected utility preference* if there is a function $u : X \to \mathbf{R}$ and a probability measure π on Ω for which $f \succeq g$ iff $\sum_{\omega} u(f(\omega))\pi_{\omega} \geq \sum_{\omega} u(g(\omega))\pi_{\omega}$.

The following result (due to T. Börgers) is related to the results in statistical decision theory known as "complete class theorems," except that no convexification via randomization is required.

Theorem 8.7 *Suppose that $\mathcal{F} \subseteq X^{\Omega}$ is finite and that $\succeq$ satisfies monotonicity. If $f^* \in \mathcal{F}$ satisfies $f^* \succ g$ for all $g \in \mathcal{F}\backslash\{f^*\}$, then there exists a subjective expected utility preference $\succeq^*$ for which $f^* \succ^* g$ for all $g \in \mathcal{F}\backslash\{f^*\}$.*

Proof. Because $f^* \in \mathcal{F}$ is strictly better than all $g \neq f^*$ where $g \in \mathcal{F}$, it follows that there must exist at least one non-null state. In the rest of the proof we ignore null states; it is easy to assign them probability zero at the end.

The proof is by induction on the size of Ω. We actually prove the slightly stronger induction hypothesis: if $\mathcal{G}$ is a finite set of acts, $\succeq$ is a preference on $\mathcal{G}$, and $g^* \in \mathcal{G}$ has the property that for all $g \in \mathcal{G}\backslash\{g^*\}$, there is $\omega \in \Omega$ for which $g^*(\omega) \succ g(\omega)$, then there is $u : X \to \mathbf{R}$ and π on $(\Omega, 2^{\Omega})$ for which $\sum_{\omega} u(g^*(\omega))\pi_{\omega} > \sum_{\omega} u(g(\omega))\pi_{\omega}$ for all $g \in \mathcal{G}\backslash\{g^*\}$. (The proof of Theorem 8.7 will then be done, as monotonicity implies that for all $f \in \mathcal{F}\backslash\{f^*\}$, there is $\omega \in \Omega$ for which $f^*(\omega) \succ f(\omega)$.)

The result is trivial if $|\Omega| = 1$. Suppose then that $|\Omega| > 1$, and that the result is true whenever the size of the state space is strictly smaller than $|\Omega|$. Let $X^* = \{x \in g^*(\Omega) : g^*(\omega) \succeq x \text{ for all } \omega \in \Omega\}$; that is, X^* is the set of worst possible outcomes occurring with g^* (recall that $g^*(\Omega)$ is finite). Consider $\Omega^* = \{\omega \in \Omega : g^*(\omega) \in X^*\}$ and $\mathcal{G}^* = \{g \in \mathcal{G} : \exists \omega \in \Omega \text{ such that for all } x \in X^*, x \succ g(\omega)\}$: Ω^* is the set of states which lead to one of the worst outcomes, and $\mathcal{G}^*$ is the set of acts which realize outcomes worse than any in X^*.

Suppose that $\Omega\backslash\Omega^* \neq \varnothing$. Clearly, $|\Omega\backslash\Omega^*| < |\Omega|$. Further, $g^* \in \mathcal{G}\backslash\mathcal{G}^*$. Finally, there is no $g \in \mathcal{G}\backslash\mathcal{G}^*$ for which $g(\omega) \succeq g^*(\omega)$ for all $\omega \in \Omega\backslash\Omega^*$. If there were, then by definition, since $g \notin \mathcal{G}^*$, g never realizes a worse outcome than an outcome in X^*; consequently, we would have $g(\omega) \succeq g^*(\omega)$ for all $\omega \in \Omega^*$ as well, so that $g(\omega) \succeq g^*(\omega)$ for all $\omega \in \Omega$, contradicting the hypothesis.

Therefore, by the induction hypothesis, there exists $u^* : X \to \mathbf{R}$ and π^* on $\Omega \setminus \Omega^*$ for which $\sum_{\omega \in \Omega\setminus\Omega^*} u^*(g^*(\omega))\pi^*_{\omega} > \sum_{\omega \in \Omega\setminus\Omega^*} u^*(g(\omega))\pi^*_{\omega}$ for all $g \in \mathcal{G} \setminus \mathcal{G}^*$. For $\delta > 0$, let $u^*_{\delta} : X \to \mathbf{R}$ be defined by $u^*(x)$ if $x \succeq x^*$ for some $x^* \in X^*$, and $u_\delta^*(x) = u^*(x) - \delta$ otherwise. Further, for $\varepsilon > 0$, define $\pi_\varepsilon^*(\{\omega\}) = \frac{\varepsilon}{|\Omega^*|}$ if $\omega \in \Omega^*$, otherwise, $\pi_\varepsilon^*(\{\omega\}) = (1-\varepsilon)\pi^*(\{\omega\})$.

For ε small enough, g^* satisfies $\sum_{\omega\in\Omega} u^*(g^*(\omega))\pi_\varepsilon^*(\{\omega\}) > \sum_{\omega\in\Omega} u^*(g(\omega))\pi_\varepsilon^*(\{\omega\})$ for all $g \in \mathcal{G}\backslash\mathcal{G}^*$. Now, for all $\delta > 0$, and for all $g \in \mathcal{G}\backslash\mathcal{G}^*$, we have $\sum_{\omega\in\Omega} u^*(g(\omega))\pi_\varepsilon^*(\{\omega\}) = \sum_{\omega\in\Omega} u_\delta^*(g(\omega))\pi_\varepsilon^*(\{\omega\})$. By choosing $\delta > 0$ large, we can ensure that for any $g \in \mathcal{G}^*$, $\sum_{\omega\in\Omega} u_\delta^*(g(\omega))\pi_\varepsilon^*(\{\omega\})$ can be made arbitrarily small. This completes the induction step.

On the other hand, if $\Omega\backslash\Omega^* = \varnothing$, we know that $g^*(\omega) \sim g^*(\omega')$ for all ω, ω'. In this case, let π be arbitrary, and let $u^* : X \to \mathbf{R}$ so that for any $x \in X^*$ and $g \in \mathcal{G}$, if $x \succ g(\omega)$, then $u^*(x) > u^*(g(\omega))$. Again by considering u_δ^* for $\delta > 0$ large, the result follows.

Theorem 8.7 establishes that, with one observation, the empirical content of monotonicity coincides with the empirical content of subjective expected utility maximization. With more than one observation, such a result does not hold in general.

Given the conclusion of Theorem 8.7, it is useful to have an understanding of the empirical content of monotonicity. We shall explain this empirical content in the case of one observation.

Proposition 8.8 *Suppose that $\mathcal{F} \subseteq X^\Omega$ is finite, and let $f^* \in \mathcal{F}$. Then there exists a monotonic preference relation $\succeq$ for which $f^* \succ g$ for all $g \in \mathcal{F} \setminus \{f^*\}$ iff $\omega_g \in \Omega$ can be chosen for all $g \in \mathcal{F}\backslash\{f^*\}$ such that the binary relation $\{(f^*(\omega_g), g(\omega_g))\}_{g\in\mathcal{F}\backslash\{f^*\}}$ on X is acyclic.*

Proof. If there exists a monotonic and rational $\succeq$, we know that for each $g \in \mathcal{F}\backslash\{f^*\}$, there is ω_g for which $f^*(\omega_g) \succ g(\omega_g)$; if not, then $g(\omega) \succeq f^*(\omega)$ for all $\omega \in \Omega$, whereby monotonicity dictates that $g \succeq f^*$, contradicting $f^* \succ g$. The relation $\{(f^*(\omega_g), g(\omega_g))\}_{g\in\mathcal{F}\backslash\{f^*\}}$ is clearly acyclic, as it is a subrelation of $\succ$ defined on X.

On the other hand, suppose that there are states $\omega_g \in \Omega$ for which the relation $\{(f^*(\omega_g), g(\omega_g))\}_{g\in\mathcal{F}\backslash\{f^*\}}$ on X is acyclic. By Szpilrajn's Theorem (Theorem 1.4) there is an extension of this binary relation to a linear order $\succeq^*$ on X. Define $\succeq'$ on $\mathcal{F}$ by $f \succeq' g$ if and only if $f(\omega) \succeq^* g(\omega)$ for all $\omega \in \Omega$: note that by definition $\succeq'$ is reflexive, transitive, and antisymmetric. Further, if $g \succeq' f^*$, then $g = f^*$. Let

$$\succeq'' = \succeq' \cup \bigcup_{g\in\mathcal{F}\backslash\{f^*\}} \{(f^*, g)\}.$$

There can be no $\langle \succeq'', \succ'' \rangle$ cycles as, by construction, if $g \succeq' f^*$, then $g = f^*$. As there are no $\succeq''$ cycles, we can use Theorem 1.5 again to extend $\succeq''$ to a preference relation $\succeq$ on $\mathcal{F}$ such that for all $g \in \mathcal{F}\backslash\{f^*\}, f^* \succ g$. The preference relation is, by construction, monotonic.

8.4 SUBJECTIVE EXPECTED UTILITY WITH AN ACT-DEPENDENT PRIOR

Many modern theories of decisions under uncertainty are predicated on an assumption that probability can depend on the act chosen; for example, the theory of maxmin expected utility supposes a utility of the form $U(f) = \sum_{\omega\in\Omega} u(f(\omega))\pi_f(\{\omega\})$, where π_f is chosen from $\arg\min_{\pi\in\Pi} \sum_{\omega\in\Omega} u(f(\omega))\pi_f(\{\omega\})$ for some $\Pi \subseteq \Delta(\Omega)$.

We end the chapter by considering the testable implications of the notion that choices are guided by an expected utility calculation in which the probability measure over states can depend on the act chosen. The notion of a dataset will correspond to that in Chapter 2, namely abstract choice.

Let X and Ω be finite sets. Say that a preference $\succeq$ on X^Ω is an *act-dependent probability representation* if there is $u : X \to \mathbf{R}$ and, for all $f \in X^\Omega$, there is $\pi_f \in \Delta(\Omega)$ such that the function $U(f) = \sum_{\omega\in\Omega} u(f(\omega))\pi_f(\{\omega\})$ represents $\succeq$.

It is quite easy to characterize act-dependent probability preferences. Say that a preference $\succeq$ satisfies *uniform monotonicity* if $f \succeq g$ whenever $f, g \in X^\Omega$ are such that for all $\omega, \omega' \in \Omega, f(\omega) \succeq g(\omega')$.[4]

Proposition 8.9 *If X^Ω is finite, a preference $\succeq$ has an act-dependent probability representation iff it satisfies uniform monotonicity.*

Proof. Suppose $\succeq$ is a preference satisfying uniform monotonicity. Since X^Ω is finite and $\succeq$ is a preference relation, there exists $U : X^\Omega \to \mathbf{R}$ which represents $\succeq$. Define $u(x) = U(x)$, where $U(x)$ is the value of U applied to the constant act taking outcome x. Note that for every f, by uniform monotonicity,

$$\min_{\omega\in\Omega} u(f(\omega)) \le U(f) \le \max_{\omega\in\Omega} u(f(\omega)).$$

Therefore, $U(f)$ can be expressed as a convex combination of $\min_{\omega\in\Omega} u(f(\omega))$ and $\max_{\omega\in\Omega} u(f(\omega))$. Choose the probability π_f so as to obtain the result of this convex combination.

The other direction is equally simple; suppose that for all $\omega, \omega' \in \Omega, f(\omega) \succeq g(\omega')$: $U(f) = \sum_{\omega\in\Omega} u(f(\omega))\pi_f(\{\omega\}) \ge \min_{\omega\in\Omega} u(f(\omega)) \ge \max_{\omega\in\Omega} u(g(\omega)) \ge \sum_{\omega\in\Omega} u(f(\omega))\pi_g(\{\omega\}) = U(g)$.

We can now speak of a choice function defined on a collection of "budgets," or sets of feasible acts Σ, each element of Σ being a subset of X^Ω. We assume that Σ is *certainty inclusive*, in the sense that for all $x, y \in X$, $\{x, y\} \in \Sigma$. Thus, choice from every pair of certain outcomes can be observed. According to this choice function, we have our standard revealed preference pair, $\langle \succeq^c, \succ^c \rangle$, as defined in Chapter 2. We introduce a new relation $\succeq'$ defined by $f \succeq' g$ if $f(\omega) \succeq^c g(\omega')$ for all $\omega, \omega' \in \Omega$.

With the property of certainty inclusiveness, the empirical content of uniform monotonicity is quite easy to describe using previous results.

[4] With a slight abuse of notation, an element of X is identified with a constant act returning that element.

Theorem 8.10 *A choice function on a certainty inclusive domain is strongly rationalized by a preference relation satisfying uniform monotonicity iff* $\langle \succeq^c \cup \succeq', \succ^c \rangle$ *is acyclic.*

Proof. Follows easily from Corollary 1.6 and Theorem 2.5.

8.5 CHAPTER REFERENCES

Allais' paradox is due to Allais (1953). He presents a classic test of the expected utility hypothesis. Typical choices in Allais' experiment directly violate the independence axiom. Theorem 8.2 appears in Fishburn (1974), among other related results. Bar-Shira (1992) also investigates the set of linear inequalities arising in this problem, and uses it to bound risk aversion in the context of monetary lotteries. Border (1992) extends this idea to a choice-theoretic approach, assuming monetary lotteries and a strictly increasing utility index. Border's approach is closer to the ideas in Afriat's Theorem. Kim (1996) also provides a generalization of this result.

The equivalence between (I) and (III) in Theorem 8.3 is essentially taken from Kubler, Selden, and Wei (2014). The equivalence between (I) and (II) in Theorem 8.3 is a version of Afriat inequalities for this problem: it appears in Green and Srivastava (1986), Varian (1983b), Varian (1988b), Bayer, Bose, Polisson, Renou, and Quah (2012), and Diewert (2012). The recent paper by Chambers, Liu, and Martinez (2014) provides a revealed preference axiom for the case of multiple goods in each state (the setup of Green and Srivastava, 1986). Green and Osband (1991) study a version of the objective probability problem in which the probability measure over states is changing, and "demand" as a function of the objective probability measure over states is observed. Park (1998) conducts a related investigation of the weighted expected utility model of Hong (1983).

The Epstein Test is due to Epstein (2000), who viewed his test as a market counterpart of Ellsberg's paradox (Ellsberg, 1961). Theorem 8.6 is based on the work of Green and Srivastava (1986), who take π as an observable, and also provide a cyclic monotonicity-style test; Kim (1991) presents related results in this direction. Bayer, Bose, Polisson, Renou, and Quah (2012) investigate related conditions which arise in the context of ambiguity models.

Axiomatizations of subjective expected utility have a long history, but the most important are due to Savage (1954) and Anscombe and Aumann (1963). The latter of these provides an axiomatization for finite states of the world. Theorem 8.6 is due to Echenique and Saito (2013), which also presents a characterization of state-dependent utility (i.e. additively separable utility across states). The equivalence between rationalizability and the version of Afriat inequalities in Theorem 8.6 is the same as obtained by Green and Srivastava (1986) (presented in Theorem 8.3 for the case of objective expected utility, only existentially quantified over the prior). The papers by Bayer, Bose, Polisson, Renou, and Quah (2012), Ahn, Choi, Gale, and Kariv (2014), and

Hey and Pace (2014) apply revealed preference tests to experimental data on uncertainty and subjective probabilities.

Polisson, Renou, and Quah (2013) also develop a system of Afriat inequalities for the model of subjective expected utility, and for other models. One important difference between the work of Polisson and Quah and other papers is that they do not require the utility over money to be concave.

Chambers, Echenique, and Saito (2015) give revealed preference axioms for translation invariant and homothetic models of choice under uncertainty, including the maxmin model of Gilboa and Schmeidler (1989), and the expected utility model when the utility function takes the specific "constant absolute risk aversion" (CARA) or "constant relative risk aversion" (CRRA) form.

The paper by Richter and Shapiro (1978) should be mentioned as well. They study the design of a set of pairwise comparisons so that, given the outcome of such a pair of comparisons, a definitive statement can be made about the agents' subjective probabilities. An example of such a statement is whether the probability of state ω_1 is at least twice that of state ω_2.

Theorem 8.7 is due to Börgers (1993), and the interpretation offered here is due to Lo (2000). It can be shown that for the primitive $\succeq$, one can choose $\succeq^*$ to have the same null states and the same ranking over constant acts as $\succeq$. This result is related to a class of results in statistical decision theory known as "complete class theorems," which establish related results when the outcome space has a convex structure due to randomization and linearity of payoffs in randomization. These results are originally due to Wald (1950, 1947a,b); Dvoretzky, Wald, and Wolfowitz (1951) use the fact that the range of a nonatomic vector measure is convex and compact to establish a similar result without any required randomization. A classic reference on these topics is Ferguson (1967).

Bossert and Suzumura (2012) establish Theorem 8.10. The question of the empirical content of preferences satisfying uniform monotonicity on domains which do not satisfy certainty inclusiveness is open, as is the same question for monotonic preferences.

Finally, there is a literature investigating the empirical content of different updating rules for probabilistic beliefs. Shmaya and Yariv (2012) characterize the empirical content of Bayes' rule, and show that it is equivalent to many apparently more general classes of updating rules.