Microeconomic Foundations I:
Choice and Competitive Markets

Student's Guide

Chapter 5: Choice under Uncertainty

Summary of the Chapter

This chapter develops and discusses the basic models that economists use to model consumer or individual choice under uncertainty. At fifty pages, it is (with Chapter 16) one of the longest chapters in this volume, and only the most intrepid reader will want to do it all in one sitting. (If you make it all the way through Section 5.2 in one sitting, you are doing great. And see below about the last bits of Section 5.2.)

The plot of the chapter, roughly, is:

- Section 5.1 introduces two frameworks for discussing choice under uncertainty, one framework with *states of nature*, a second with *(objective) probability distributions over prizes*. It also introduces the two *expected utility* representations that are sought in the rest of the chapter: using the given probabilities in the second context; and employing probabilities over the states of nature that are derived from the decision maker's preferences in the first context. (Other names you will encounter in the literature are: the *von Neumann–Morgenstern expected utility model* or *representation* when the probabilities are given "objectively"; and the *subjective expected utility model* or the *Savage model* when the probabilities over states of nature are derived from the decision maker's preferences.)
- Section 5.2 does the heavy lifting for objective probabilities and expected utility, employing a more general result known as the *Mixture-Space Theorem*. At first everything is done for probability distributions that have finite support, so-called *simple probability distributions*, on a given space of prizes. And then the section discusses how one extends from simple to more complex probability distributions, and

why (in theory) this often comes at the price of bounded utility functions. This final discussion will be tough sledding for readers whose mathematical background doesn't include weak convergence of probability measures; if that is you, you can safely go on to Section 5.3 without a full understanding of this part of Section 5.2.

- Section 5.3 does the somewhat less-heavy lifting for subjective expected utility. While the classic derivation is due to L. J. Savage (1954), his treatment is too long and complex for this book, and we employ instead the very clever development of Anscombe and Aumann (1963). After all the hard work of Section 5.2, the mathematical derivations in this section seem quite simple. But this hides what is really going on, and an important part of the section is a discussion of how, in the Anscombe and Aumann treatment, the *formulation* of the problem contains a critical assumption. To appreciate fully this point, I urge you to do Problems 5.4 and 5.5 and, to the extent possible, discuss the solutions to those problems with your peers/in class.
- Having finished with the mathematical development, Section 5.4 discusses the point of view held by many economic theorists (but not by me) that differences in subjective probability assessment can only result from differences in information, the so-called *common-prior assumption.*
- The chapter concludes in Section 5.5 with a discussion of empirical and theoretical reasons to doubt the validity of the models developed in this chapter.

Solutions to Starred Problems

■ 5.1. (a) Gamble 1 has expected utility $(1.0)[\sqrt{10{,}000}] = 100$. Gamble 2 has expected utility $(1/3)[\sqrt{3600}] + (2/3)[\sqrt{14{,}400}] = 20 + 80 = 100$. Gamble 3 has expected utility $(1/5)[\sqrt{0}] + (1/5)[\sqrt{10{,}000}] + (2/5)[\sqrt{22{,}500}] = 0 + 20 + 90 = 110$. So the consumer will select gamble 3.

(b) The first act evaluates as $\sqrt{4} + 0.6 \min\{3, 12\} + 0.4(7 + 7) = 2 + 1.8 + 5.6 = 9.4$. The second evaluates as $\sqrt{16} + 0.6 \min\{2, 20\} + 0.4(10 + 0) = 4 + 1.2 + 4 = 9.2$. The third evaluates as $\sqrt{25} + 0.6 \min\{5, 5\} + 0.4(5 + 5) = 5 + 3 + 4 = 12$. This consumer will choose the third act.

(c) The first act evaluates as $(0.5)(4)^{0.25} + (0.3)(36)^{0.25} + (0.2)(49)^{0.25} = 1.97$. The act evaluates as $(0.5)(16)^{0.25} + (0.3)(40)^{0.25} + (0.2)(0)^{0.25} = 1.754$. The act evaluates as $(0.5)(25)^{0.25} + (0.3)(25)^{0.25} + (0.2)(25)^{0.25} = 2.236$. So the third act is chosen.

(d) The first horse-race lottery evaluates as 13.92, and the second at 12. So the first is chosen.

(e) The first horse-race lottery evaluates at 2.214 and the second as 2.236, so the second is chosen.

■ 5.3. The problem is to prove that a through c in the Mixture-Space Theorem (Proposition 5.4) imply the representation d and e, without the extra assumption enlisted in the proof in the text.

Note that the lemmas in the proof do not use this extra assumption, so we have them all available.

So assume we have a mixture space Z and a preference relation $\succeq$ on Z that satisfies a through c. Either $z \sim z'$ for all z and z' from Z, or $z \succ z'$ for some pair z and z'. In the first case, any constant function u satisfies d and e, so we can proceed with the second case. Fix a pair, denoted $\overline{z}$ and $\underline{z}$, such that $\overline{z} \succ \underline{z}$.

Now choose any $z \in Z$ and define $u(z)$ as follows:

Case A: If $\overline{z} \succeq z \succeq \underline{z}$, let a be the unique number between zero and one such that $z \sim a\overline{z} + (1-a)\underline{z}$, and set $u(z) = a$.

Case B: If $z \succ \overline{z}$, let a be the unique number strictly between zero and one such that $\overline{z} \sim az + (1-a)\underline{z}$, and let $u(z) = 1/a$.

Case C: If $\underline{z} \succ z$, let a be the unique number strictly between zero and one such that $\underline{z} \sim a\overline{z} + (1-a)z$, and let $u(z) = a/(a-1)$.

Why are these scalars a all unique? In cases B and C, why are they strictly between 0 and 1? We note first that in Case A, $\overline{z} \succ \underline{z}$ by assumption, in case B, $z \succ \overline{z} \succ \underline{z}$, and in case C, $\overline{z} \succ \underline{z} \succ z$; then apply the lemmas and, in particular, Lemma 5.9.

We need to show that $u(z)$ so defined represents $\succ$ and is linear in convex combinations. In other words, if z and z' are both from Z and if $\alpha \in [0,1]$, then $z \succeq z'$ if and only if $u(z) \geq u(z')$, and $u(\alpha z + (1-\alpha)z') = \alpha u(z) + (1-\alpha)u(z')$.

So fix for the remainder of this solution some z and z' from Z and $\alpha \in [0,1]$.

Let z_1 be the $\succeq$-best of the threesome z, z', and $\overline{z}$ (if there is a tie for best, choose one of the best however you prefer) and let z_2 be the $\succeq$-worst of the threesome z, z' and $\underline{z}$ (again choosing one of the worst if there is a tie). Since $\overline{z} \succ \underline{z}$, $z_1 \succeq \overline{z}$, and $\underline{z} \succeq z_2$, this ensures that $z_1 \succ z_2$. Define

$$Z' = \{z \in Z : z_1 \succeq z \succeq z_2\}.$$

By the manner in which z_1 and z_2 where chosen, z, z', $\overline{z}$, and $\underline{z}$ are all in Z'. Moreover, I assert that: *If $\hat{z}$ and $\check{z}$ are both in Z' and if $a \in [0,1]$, then $a\hat{z} + (1-a)\check{z} \in Z'$.* To see why, take any $\hat{z}$ and $\check{z}$ from Z'. By completeness of $\succeq$, we know that either $\hat{z} \succeq \check{z}$ or $\check{z} \succeq \hat{z}$. Without loss of generality, suppose that $\hat{z} \succeq \check{z}$. Then Lemma 5.7 tells us that $\hat{z} \succeq a\hat{z} + (1-a)\check{z} \succeq \check{z}$ for all $a \in [0,1]$. But $z_1 \succeq \hat{z}$ and $\check{z} \succeq z_2$, so transitivity of $\succeq$ gives us $z_1 \succeq a\hat{z} + (1-a)\check{z} \succeq z_2$ for all a.

This implies that Z' is a mixture space: It is closed under the taking of convex combinations. And since a, b, and c from Proposition 5.3 hold on all of Z and are "for all" statements, they hold for all $z \in Z'$. So Z' is a mixture space satisfying a, b, and c. What is more, Z' contains a $\succeq$-best and a $\succeq$-worst element, namely z_1 and z_2, respectively. So the proof of Proposition 5.3 given in the text works perfectly well for Z': There exists a function $v : Z' \to R$ that represents $\succeq$ on Z' and that is linear in convex combinations.

Moreover, Proposition 5.3 tells us that this is true as well of any $v' : Z' \to R$ that has $v' = Av + B$, for any strictly positive scalar A and any scalar B.

We know that $\overline{z} \succ \underline{z}$, and of course both are elements of Z', so $v(\overline{z}) > v(\underline{z})$. So let

$$A = \frac{1}{v(\overline{z}) - v(\underline{z})} \quad \text{and} \quad B = \frac{-v(\underline{z})}{v(\overline{z}) - v(\underline{z})},$$

and let $v' = Av + B$ for this specific choice of $A > 0$ and B. Simple algebra gives $v'(\underline{z}) = Av(\underline{z}) + B = v(\underline{z})/(v(\overline{z}) - v(\underline{z})) - v(\underline{z})/(v(\overline{z}) - v(\underline{z})) = 0$, and $v'(\overline{z}) = Av(\overline{z}) + B = v(\overline{z})/(v(\overline{z}) - v(\underline{z})) - v(\underline{z})/(v(\overline{z}) - v(\underline{z})) = 1$.

Moreover, I asset that *for any* $\hat{z} \in Z'$, $v'(\hat{z}) = u(\hat{z})$. That is, v' coincides with the originally defined u on all of Z'. We go back to the three cases at the start of the problem, for $\hat{z}$ this time:

Case A: If $\overline{z} \succeq \hat{z} \succeq \underline{z}$, $u(a)$ is the unique number a such that $\hat{z} \sim a\overline{z} + (1-a)\underline{z}$. But if $\hat{z} \sim a\overline{z} + (1-a)\underline{z}$, $v'(\hat{z}) = v'(a\overline{z} + (1-a)\underline{z}) = av'(\overline{z}) + (1-a)v'(\underline{z}) = a$, since $v'(\overline{z}) = 1$ and $v'(\underline{z}) = 0$.

Case B: If $\hat{z} \succ \overline{z}$, $u(\hat{z})$ is $1/a$, where a is the unique number strictly between zero and one such that $\overline{z} \sim a\hat{z} + (1-a)\underline{z}$. But then $v'(\overline{z}) = v'(a\hat{z} + (1-a)\underline{z}) = av'(\hat{z}) + (1-a)v'(\underline{z})$, or $1 = av'(\hat{z})$, or $v'(\hat{z}) = 1/a$.

Case C: If $\underline{z} \succ \hat{z}$, $u(\hat{z}) = a/(a-1)$, where a is the unique number strictly between zero and one such that $\underline{z} \sim a\overline{z} + (1-a)\hat{z}$. But then $v'(\underline{z}) = v'(a\overline{z} + (1-a)\hat{z}) = av'(\overline{z}) + (1-a)v'(\hat{z})$, giving us the equation $0 = a + (1-a)v'(\hat{z})$, which solves as $v'(\hat{z}) = a/(a-1)$.

Please note that in all three cases, I am repeatedly using the fact that for any two elements of Z', their convex combinations are also in Z'.

This (practically) finishes the proof. Since v' coincides with u on Z' and v' represents $\succeq$ and is linear in convex combinations on Z', we know that u represents $\succeq$ and is linear in convex combinations on Z'. And our originally (arbitrarily) chosen z and z' are both in Z', so we know that

$$z \succeq z' \text{ if and only if } u(z) \geq u(z'), \text{ and } u(\alpha z + (1-a)z') = \alpha u(z) + (1-\alpha)u(z').$$

That's what we needed to show.

■ 5.6. This solution will be completely unintelligible if you don't have the textbook open in front of you, as I'm not going to repeat the notation employed.

a. H is a mixture space, and the Mixture-Space Theorem applies. If $\succeq$ satisfies the three mixture-space axioms, there exists a function $F : H \to R$ that represents $\succeq$ and that is linear in convex combinations. (The other parts of the Mixture-Space Theorem, concerning the converse and the extent to which the representing function F is "unique," are also all true, of course.)

b. The key is to employ the "trick" of Anscombe and Aumann, namely

$$\frac{1}{n}[\ldots, \pi_{ij} \text{ on } A_{ij}, \ldots] + \frac{n-1}{n}[\underline{x} \text{ on } S] = \sum_{i=1}^{n} \frac{1}{n}[\pi_{i1} \text{ on } A_{i1}, \ldots, \pi_{im} \text{ on } A_{im}, \underline{x} \text{ on } A_i^C].$$

Therefore,

$$\frac{1}{n}F([\ldots, \pi_{ij} \text{ on } A_{ij}, \ldots]) + \frac{n-1}{n}F([\underline{x} \text{ on } S]) = \sum_{i=1}^{n} \frac{1}{n}F([\pi_{i1} \text{ on } A_{i1}, \ldots, \pi_{im} \text{ on } A_{im}, \underline{x} \text{ on } A_i^C]).$$

Since $F([\underline{x} \text{ on } S]) = 0$, that term disappears, and if you multiply through by n and recall the definition of F_A for any set A, we have part b.

c. In this context, Savage's Independence Axiom would say: Suppose A and B, both from $\mathcal{A}$, partition S. For $h_A, h'_A \in H_A$ and $h_B, h'_B \in H_B$,

$$[h_A \text{ on } A, h_B \text{ on } B] \succeq [h'_A \text{ on } A, h_B \text{ on } B] \text{ if and only if}$$
$$[h_A \text{ on } A, h'_B \text{ on } B] \succeq [h'_A \text{ on } A, h'_B \text{ on } B].$$

That is, if we compare two h's that agree on B, what is important to the comparison is how they disagree where they disagree. If we change how they agree where they agree, we don't change how they compare.

With part b, showing this is trivial:

$$\begin{aligned}
&[h_A \text{ on } A, h_B \text{ on } B] \succeq [h'_A \text{ on } A, h_B \text{ on } B] \text{ if and only if}\\
&F([h_A \text{ on } A, h_B \text{ on } B]) \geq F([h'_A \text{ on } A, h_B \text{ on } B]) \text{ if and only if}\\
&F_A(h_A) + F_B(h_B) \geq F_A(h'_A) + F_B(h_B) \text{ if and only if}\\
&F_A(h_A) + F_B(h'_B) \geq F_A(h'_A) + F_B(h'_B) \text{ if and only if}\\
&F([h_A \text{ on } A, h'_B \text{ on } B]) \geq F([h'_A \text{ on } A, h'_B \text{ on } B]) \text{ if and only if}\\
&[h_A \text{ on } A, h'_B \text{ on } B] \succeq [h'_A \text{ on } A, h'_B \text{ on } B].
\end{aligned}$$

d. Fix some $A \in \mathcal{A}$ and look at $\succeq$ restricted to h that take the form $[\pi \text{ on } A, \underline{x} \text{ on } A^C]$, for $\pi \in \Pi$. If $\succeq$ satisfies the three mixture-space axioms overall, it satisfies them on this restricted set as well, and so the von Neumann-Morgenstern Expected-Utility Theorem (Proposition 5.3) says that we have an expected utility representation, which also gives a mixture-space representation. Since F is a mixture-space representation, already, it can be used to provide the expected utility representation, where the utility of the prize x is the representing function evaluated at the lottery that gives x with certainty. Which is precisely how U_A is defined in this part of the problem. That is,

$$F_A([p \text{ on } A]) = F([p \text{ on } A, \underline{x} \text{ on } A^C]) = \sum_x U_A(x)p(x),$$

where the sum is over x in the support of p. Now apply part b.

e. For $\mathbf{p}$ as defined in the statement of the problem, $\mathbf{p}(S) = F([\overline{x} \text{ on } S]) = 1$. And if A and B are disjoint sets, then

$$\frac{1}{2}[\overline{x} \text{ on } A \cup B, \underline{x} \text{ on } (A \cup B)^C] + \frac{1}{2}\underline{x} = \frac{1}{2}[\overline{x} \text{ on } A, \underline{x} \text{ on } A^C] + \frac{1}{2}[\overline{x} \text{ on } B, \underline{x} \text{ on } B^C].$$

Apply the Mixture-Space Theorem to get

$$\frac{1}{2}F([\overline{x} \text{ on } A\cup B, \underline{x} \text{ on } (A\cup B)^C])+\frac{1}{2}F(\underline{x}) = \frac{1}{2}F([\overline{x} \text{ on } A, \underline{x} \text{ on } A^C])+\frac{1}{2}F([\overline{x} \text{ on } B, \underline{x} \text{ on } B^C]).$$

Recall that $F(\underline{x}) = 0$, so that term drops. Multiply on both sides by 2, and this is

$$\mathbf{p}(A \cup B) = \mathbf{p}(A) + \mathbf{p}(B).$$

So we have finite additivity.

We haven't used the assumption that $\overline{x} \succeq_A \underline{x}$ yet, so where does that come in? We need that to ensure that $\mathbf{p}(A) = F([\overline{x} \text{ on } A, \underline{x} \text{ on } A^C])$ is nonnegative (since we know

$F(\underline{x}) = 0$); otherwise, we produce not a finitely additive probability but a finitely additive signed measure on $(S, \mathcal{A})$.

Finally, we need to show that $\mathbf{p}(A) = 0$ if and only if A is null. But if A is null, then $\overline{x} \sim_A \underline{x}$, hence $F([\overline{x} \text{ on } A, \underline{x} \text{ on } A^C]) = F([\underline{x} \text{ on } A, \underline{x} \text{ on } A^C]) = F([\underline{x} \text{ on } S]) = 0$. And if A is not null, then $\overline{x} \succ_A \underline{x}$, and $\mathbf{p}(A) = F([\overline{x} \text{ on } A, \underline{x} \text{ on } A^C]) > F([\underline{x} \text{ on } A, \underline{x} \text{ on } A^C]) = F([\underline{x} \text{ on } S]) = 0$.

f. Given everything that has come before (and, in particular, part d), we are done once we show that U_A can be written $\mathbf{p}(A)U$ for some single function $U : X \to R$. For the function U, take $U(x) = F([x \text{ on } S])$. Then our normalization of F means that $U(\overline{x}) = 1$ and $U(\underline{x}) = 0$.

Take any $A \in \mathcal{A}$. If A is null, then $\pi \sim_A \pi'$ for all $\pi, \pi' \in \Pi$ and, in particular, $x \sim_A \underline{x}$ for all $x \in X$. Therefore, $U_A(x) = U_A(\underline{x}) = 0$ for all x. But then $\mathbf{p}(A)U(\cdot) = 0 = U_A(\cdot)$, since $\mathbf{p}(A) = 0$, and we are done.

And if A is not null, then $\succeq_A$ gives the same preference ordering on Π as does $\succeq$. Hence U_A must be a positive affine translate of U. We've fixed $U(\underline{x}) = U_A(\underline{x}) = 0$, hence $U_A(\cdot) = aU(\cdot)$ for some positive constant a. But we know that $U(\overline{x}) = 1$ and $U_A(\overline{x}) = \mathbf{p}(A)$, so this positive constant a must be $\mathbf{p}(A)$, completing the proof.

■ 5.7 or, **Concerning Proposition 5.10**:[1]

The topology of weak convergence (on the space P of Borel probability measures on $X = R^k$ or R^k_+) has the following "definition": A sequence $\{p_n\}$ out of P converges to $p \in P$ if

$$\lim_{n\to\infty} \int_X f(x)p_n(dx) = \int_X f(x)p(dx) \text{ for all bounded and continuous real-valued functions } f : X \to R.$$

The scare-quotes around *definition* are there because to define a topology, one ought to specify an appropriate collection of open sets or, at least, a neighborhood base for the topology; but given the definition of sequential convergence I've provided, readers who know about such things can probably produce for themselves a neighborhood base for this topology. We won't need to know this, but it turns out that this topology

[1] In early printings of the book, the discussion beginning on page 98 and leading to Proposition 5.10 referred to the *weak topology on P*. This is, at best, ambiguous; mathematically sophisticated readers are entitled to object that, in the most natural interpretation, it is wrong. What I wanted was the *the topology of weak convergence of probability measures*, which is a weak* topology and which I define in the text immediately following. Princeton University Press allowed me to correct typos in the book for later printings, and I fixed the language around Proposition 5.10 and added Problem 5.7, which asks the reader to prove the Proposition (now stated less ambiguously). If your copy of the book lacks Problem 5.7, you have an early printing, not corrected for typos, in which case the discussion to follow should help clarify the ambiguous text. If your copy of the book has Problem 5.7, then here is its solution.

on P is metrizable—look for references to the Prohorov metric in the literature. The facts we do need to know are:

- For every $p \in P$, we can construct a sequence $\{p_n\}$ consisting of *simple* probabilities on X that has the limit p.

- Suppose $p \in P$ has countable support $\{x_1, x_2, \ldots\}$, and we define a sequence of simple probabilites $\{p_m\}$, where p_m has support $\{x_0, \ldots, x_m\}$ (for arbitrary x_0) and $p_m(x_n) = p(x_n)$ for $n = 1, \ldots, m$ and $p_m(x_0) = \sum_{n=m=1}^{\infty} p(x_n)$. In words, p_m "truncates" the support of p at element #m, assigning any left-over probability to x_0. Then $\lim_m p_m = p$ in the topology of weak convergence. (I'm not going to give you a proof of this, but since I've characterized what it takes for a sequence to converge, you shouldn't have huge difficulty in proving this on your own.)

- Suppose $\{p_n\}$ is any sequence out of P with limit p in this topology; then for all $a \in [0, 1]$ and $p' \in P$, $\lim_n ap_n + (1 - a)p' = ap + (1 - a)p'$. (This is even easier to prove from first principles, if your first principles are the definition of sequential convergence given above.)

Using these three facts,[2] here is a proof of the Proposition:

Step 1 is to invoke the mixture-space theorem: There exists some function $u : P \to R$ that represents $\succeq$ and that is linear in convex combinations. We now aim to show that this u takes the form of expected utility and where, moreover, the utility function U is just $U(x) := u(\delta_x)$. (Of course, if u takes the form of expected utility with U, then $U(x)$ must equal $u(\delta_x)$, since that is the expected utility of the lottery δ_x.) It should be noted that for all $p \in P_S$, the induction argument used in the chapter tells us already that $u(p) = \sum_{x \in \text{supp}(p)} u(\delta_x)p(x) = \sum_{x \in \text{supp}(p)} U(x)p(x)$.

Step 2 disposes of a trivial case. Suppose $\delta_x \sim \delta_y$ for all x, $y \in X$. Then (I assert) $p \sim q$ for all $p, q \in P$, so any constant function does the trick (and, of course, $U(x) = u(\delta_x)$ is one such constant function). To see why my assertion is true, suppose $\delta_x \sim \delta_y$ for all x and y. Then $U(x) = U(y)$ for all x and y; that is, U is a constant function. But then we know (from the representation on P_S or directly by induction) that $p \sim q$ for all simple probabilities p and q. Suppose $p \succ q$ for some (non-simple) p and q from P. Let $\{p_n\}$ be a sequence of simple probabilities with limit p, and let $\{q_n\}$ be a sequence of simple probabilities with limit q; continuity and $p \succ q$ would imply that for n large enough (for all large enough n), $p_n \succ q_n$. But we know that for simple probabilities, universal indifference must hold, a contradiction. So in this case, every Borel lottery p is indifferent to every other one, and we are finished.

Hence I can assume for the remainder of the proof that there exist $\overline{x}$ and $\underline{x}$ from X such that $\delta_{\overline{x}} \succ \delta_{\underline{x}}$.

[2] And, if your copy of the book says "the weak topology," changing this to "the topology of weak convergence,"

Step 3: The function U is continuous. Suppose not. Then there is some prize x and sequence of prizes $\{x_n\}$ with limit x such that $U(x) \neq \lim_n U(x_n)$. (By looking along subsequences, I can assume that $\lim_n U(x_n)$ exists, although it may be $+\infty$ or $-\infty$.) Again I need to look at cases, one case where the limit exceeds $U(x)$; the other where it is less. I'll assume $\lim_n U(x_n) > U(x)$ and let you handle the other case. Then we know that for some $\epsilon > 0$ and for all sufficiently large n, $U(x_n) \geq U(x) + \epsilon$. Choose any one of the x_n's for which this is true and call it x', so that $U(x') \geq U(x) + \epsilon$. This means that $\delta_{x'} \succ \delta_x$. We can find an $a \in (0,1)$ such that $aU(x)+(1-a)U(x') = U(x)+\epsilon/2$; this is a simple application of the intermediate value theorem. Of course, this implies that $a\delta_x + (1-a)\delta_{x'} \succ \delta_x$. But then for all sufficient large n,

$$U(x_n) \geq U(x) + \epsilon > U(x) + \epsilon/2 = aU(x) + (1-a)U(x') = u(a\delta_x + (1-a)\delta_{x'}),$$

so $\delta_{x_n} \succ a\delta_x + (1-a)\delta_{x'}$. By continuity, since δ_{x_n} has limit δ_x in the topology of weak convergence (since x_n has limit x), this implies that $\delta_x \succeq a\delta_x + (1-a)\delta_{x'}$, which contradicts $a\delta_x + (1-a)\delta_{x'} \succ \delta_x$.

Step 4: The function U is bounded. Again, suppose not. Then it is either unbounded above or below (or both). I'll do the case where U is unbounded above, and you can deal with unbounded below. Since U is unbounded above, we can find for $n = 1, 2, \ldots$ a prize x_n such that $U(x_n) > 1/2^n$. Let p be the (countable-support, non-simple) lottery that has prize x_n with probability $1/2^n$, and let p_n be the (finite support, hence simple) lottery that has prize x_m with probability $1/2^m$ for $m \leq n$ and some fixed prize x_0 with the remaining probability $1/2^n$. It isn't hard to see that $\lim_n p_n = p$ in the topology of weak convergence.

It is here that we employ $\delta_{\overline{x}}$ and $\delta_{\underline{x}}$. Since $\delta_{\overline{x}} \succ \delta_{\underline{x}}$, the third (continuity) mixture-space axiom tells us that for some $a > 0$, $\delta_{\overline{x}} \succ ap + (1-a)\delta_{\underline{x}}$. Again relying on properties of the topology of weak convergence, we observe that $\lim_n(ap_n+(1-a)\delta_{\underline{x}}) = ap+(1-a)\delta_{\underline{x}}$, hence by continuity of preference, for all sufficient large n, $\delta_{\overline{x}} \succ ap_n + (1-a)\delta_{\underline{x}}$. But $ap_n + (1-a)\delta_{\underline{x}}$ is a simple probability, and we (therefore) know that its utility is its expected utility under U. And for any $a > 0$, the way p and the p_n were constructed implies that the expected utility of $ap_n + (1-a)\delta_{\underline{x}}$ increases without bound, as long as $a > 0$. So for all sufficiently large n, we have $ap_n + (1-a)\delta_{\underline{x}} \succ \delta_{\overline{x}}$. This is a contradiction; U must be bounded.

Step 5. The punchline. Since U is bounded and continuous, $\int_X U(x)p(dx)$ is well defined for all Borel p. We merely need to confirm that $u(p)$ is this integral for all Borel p, and we're done. I'm going to give a slightly clumsy proof of this, but one that doesn't need any special knowledge of the theory of integration, besides what we've assumed so far.

Take any p. If $u(p) \neq \int_X U(x)p(dx)$, then either $u(p) >$ the integral or less. I'll deal with the case that it is greater; you can supply the details for the case where it is less. Let $u(p) - \int_X U(x)p(dx) = \epsilon > 0$.

Take some sequence of simple probabilities $\{p_n\}$ with limit p in the topology of weak convergence. Since U is bounded and continuous, we know that $\int_X U(x)p(dx) = \lim_n \int_X U(x)p_n(dx)$; that's implied by convergence in this topology. Hence, for some N and all $n > N$,

$$u(p_n) = \int_X U(x)p_n(dx) < \int_X U(x)p(x) + \epsilon/3.$$

Of course, this means that, for all $n > N$, $u(p_n) + 2\epsilon/3 < u(p)$; hence $p \succ p_n$ for all $n > N$. Re-lable p_{N+1} as p', and let $a \in (0,1)$ be such that $u(ap + (1-a)p') = au(p) + (1-a)u(p') = \int_X U(x)p(dx) + 2\epsilon/3$. Then for all $n > N$, $p_n \prec ap + (1-a)p'$. Since $\lim_n p_n = p$, this implies $p \preceq ap + (1-a)p'$, which contradicts $p \succ p'$. ■