Crash Course on Subjective Expected Utility

The Anscombe–Aumann Approach

Lasha Chochua

ISET, 2026

Table of contents

Standard references for this material are Anscombe and Aumann (1963), Gilboa (2009) Chapters 4–5, and Kreps (1988). The present note follows the treatment in those sources, with proof structure modelled on the companion note on vNM expected utility.

1 From Objective to Subjective Probabilities

The von Neumann–Morgenstern theorem takes probabilities as *given*. The set of lotteries ΔX consists of probability distributions over outcomes X, and the probabilities p_i are part of the description of the objects being ranked. Every decision maker faces the same set of lotteries: they can only differ in the vNM utility index v.

This assumption is restrictive. Most uncertainty in economics is not of this kind. When a firm decides whether to enter a market, when an investor chooses a portfolio, or when a policymaker evaluates an intervention, the relevant probabilities are *not* objectively given – they reflect the decision maker's own beliefs about the world. Two agents facing the same situation may disagree about the likelihood of each outcome, and neither can be said to be objectively wrong.

To accommodate this, we need a richer model in which both beliefs (a probability distribution over states of the world) *and* tastes (a utility index over consequences) are characteristics of the individual, derived jointly from her preferences. This is the programme of *subjective expected utility* (SEU). Anscombe and Aumann (1963) provide the cleanest axiomatisation, by embedding an objective randomisation device (the "roulette") inside a subjectively uncertain environment (the "horse race"). The result is a two-stage structure that separates the two sources of uncertainty and allows each to be identified from preferences.

2 Setup

2.1 Primitives

States. Let $\Omega = \{1, 2, \dots, S\}$ be a finite set of *states of the world*, with generic element $s \in \Omega$. States represent mutually exclusive, exhaustive descriptions of the world. The decision maker does not know which state obtains.

Outcomes. Let $X = \{x_1, \dots, x_n\}$ be a finite set of *outcomes*, with generic element $x \in X$.

Roulette lotteries. ΔX is the set of all probability distributions over X: elements $\pi \in \Delta X$ with $\pi_i \geq 0$ and $\sum_i \pi_i = 1$. These are *objective* lotteries – their probabilities are publicly known and agreed upon, as if determined by a physical roulette wheel.

Acts. An *Anscombe–Aumann act* is a function $h\colon \Omega \to \Delta X$. The space of all acts is $H = (\Delta X)^{\Omega}$. An act assigns to each state a roulette lottery over consequences. Write $h_s = h(s) \in \Delta X$ for the lottery assigned to state s, and $h_s(x) \in [0, 1]$ for the probability of outcome x under that lottery, so $h_s(x) = \Pr(x \mid s, h)$.

Primitive. A binary relation $\succsim$ on H (preference over acts), with $\succ$ and $\sim$ defined as usual.

2.2 The Horse Race and Roulette Interpretation

The two-stage structure of H has a natural interpretation introduced by Anscombe and Aumann. Think of Ω as the field of horses in a race: state s is the event that horse s wins. The decision maker *subjectively* assesses each horse's chance of winning – different decision makers may disagree. Think of ΔX as the outcomes of a roulette spin: probabilities are *objectively* determined, the same for everyone.

An act $h \in H$ is then a voucher that says: if horse s wins the race, you will be entered into roulette lottery h_s over prizes X. First the horses run (subjective uncertainty); then, conditional on the winner, the roulette is spun (objective uncertainty).

This separation is what makes Anscombe–Aumann cleaner than Savage (1954). Because the second stage is an objective lottery, the vNM theorem can be applied state by state to evaluate h_s. The harder problem – identifying subjective beliefs over Ω – is then handled by a separate axiom (state-independence) imposed on top.

2.3 Three Descriptions of H

An act $h \in H$ can be described in three equivalent ways.

As a function. $h\colon \Omega \to \Delta X$, directly. For example with $\Omega = \{s_1, s_2, s_3\}$ and $X = \{x_1, x_2, x_3\}$:

$$h(s_1) = (0.3,\ 0.2,\ 0.5), \quad h(s_2) = (0.4,\ 0.6,\ 0), \quad h(s_3) = (0,\ 1,\ 0).$$

As a compound lottery. A two-stage tree: the first stage (subjective) picks $s \in \Omega$; the second stage (objective) picks $x \in X$ according to h_s.

As a matrix. H is the set of weakly positive $S \times n$ matrices with each row summing to 1:

$$h = \begin{pmatrix} 0.3 & 0.2 & 0.5 \\ 0.4 & 0.6 & 0 \\ 0 & 1 & 0 \end{pmatrix},$$

where rows are states and columns are outcomes. Entry $h_{s,x}$ is the probability of x in state s: $h_s(x) = h_{s,x}$.

2.4 Mixing Acts and Constant Acts

Because ΔX is convex, H inherits a linear structure. For $f, g \in H$ and $\alpha \in [0,1]$, the *mixture* $\alpha f + (1-\alpha)g$ is the act defined pointwise:

$$[\alpha f + (1-\alpha)g](s) = \alpha f(s) + (1-\alpha)g(s) \quad \forall s \in \Omega.$$

Mixing operates on the roulette lotteries state by state. It does not mix the horses – those are resolved by nature. This pointwise mixing operation is what allows the Independence and Archimedean axioms to be stated for acts.

A *constant act* is one that assigns the same lottery regardless of the state: $f \in H$ is constant if $f(s) = f(s')$ for all $s, s' \in \Omega$. The set of constant acts H_c is isomorphic to ΔX, and we identify the two. In particular, any $\pi \in \Delta X$ can be treated as the constant act that returns π in every state. Similarly, any outcome $x \in X$ corresponds to the constant act returning the degenerate lottery δ_x in every state. Thus $X \subset \Delta X \subset H$, and A–A acts strictly generalise the vNM setting.

3 The Axioms

The first four axioms are exactly the vNM axioms, now applied to the richer domain H of acts rather than to ΔX directly. The fifth is new.

Completeness and Transitivity. $\succsim$ is a complete and transitive binary relation on H.

Independence. For all $f, g, k \in H$ and all $\alpha \in (0, 1)$:

$$f \succsim g \iff \alpha f + (1-\alpha)k \succsim \alpha g + (1-\alpha)k.$$

Archimedean (Continuity). For all $f, g, k \in H$ with $f \succ g \succ k$, there exist $\alpha, \beta \in (0, 1)$ such that:

$$\alpha f + (1-\alpha)k \succ g \succ \beta f + (1-\beta)k.$$

State-Independence. For all non-null states $s, t \in \Omega$, all acts $h, g \in H$, and all lotteries $\pi, \rho \in \Delta X$:

$$(h_{-s}, \pi) \succsim (h_{-s}, \rho) \implies (g_{-t}, \pi) \succsim (g_{-t}, \rho).$$

Here (h_{-s}, π) denotes the act that coincides with h in every state except s, where it assigns lottery π.

Before discussing what each axiom does, we need the concept of a null state.

A state $s \in \Omega$ is *null* if $(h_{-s}, \pi) \sim (h_{-s}, \rho)$ for all $h \in H$ and all $\pi, \rho \in \Delta X$. A state is *non-null* if it is not null: there exist h, π, ρ such that $(h_{-s}, \pi) \succ (h_{-s}, \rho)$.

A null state is one the decision maker assigns probability zero to – changing what happens there never affects preferences.

3.1 What Each Axiom Does

Completeness and Transitivity are the bare minimum for any representation: they guarantee a complete, transitive ranking of acts, without which no numerical representation exists.

Independence over acts says that mixing two acts f and g with a common third act k should not reverse their ranking. As in vNM, this is the critical axiom for the *scale type* of the representation. Applied to constant acts (elements of $\Delta X \subset H$), it recovers the vNM independence axiom, which forces utility over roulette lotteries to be linear in probabilities. Applied to non-constant acts, it does further work: it forces the representation to be *additive* across states.

Archimedean rules out infinitely good or bad acts. Combined with Independence, it ensures continuous mixture-space structure on H, which is what allows the Mixture Space Theorem to apply.

Together, Completeness, Transitivity, Independence, and Archimedean are exactly the conditions of the Mixture Space Theorem applied to H. This gives a *state-dependent* additive representation – but not yet SEU, because utility may vary across states.

State-Independence is the additional axiom that separates A–A from the state-dependent intermediate result. It requires that the ranking of roulette lotteries π versus ρ does not depend on the state in which they are received, provided that state is non-null. This axiom forces the vNM index over consequences to be the same function v in every state, which in turn allows subjective probabilities to be identified as separate objects.

4 The Intermediate Result: State-Dependent Utility

Before stating the main theorem it is instructive to see how far Axioms 1–4 alone take us, without state-independence.

Since $H = (\Delta X)^\Omega$ is a convex subset of $(\mathbb{R}^n)^\Omega$, the Mixture Space Theorem applies directly. The theorem says: if $\succsim$ on a convex set M is complete, transitive, independent, and Archimedean, then $\succsim$ has an affine utility representation. An affine function on $H \subset \mathbb{R}^{S\times n}$ takes the form of a linear function on each state's lottery.

State-Dependent Representation. The preference relation $\succsim$ on H satisfies Axioms 1–4 if and only if there exist functions $v_1, \dots, v_S \colon X \to \mathbb{R}$ such that

$$U(h) = \sum_{s\in\Omega} \sum_{x\in X} v_s(x)\, h_s(x) \tag{1}$$

is a utility representation of $\succsim$.

The key feature of (1) is that the utility index v_s is *state-dependent*: the value of outcome x may differ depending on which state of the world obtains. This means one cannot separate probabilities from utilities in (1). Write $v_s(x) = \mu(s) \cdot v(x)$ for some $\mu(s) > 0$ and $v(x)$: this factorisation is one valid representation, but so is $v_s(x) = \mu'(s) \cdot v'(x)$ for infinitely many other pairs (μ', v'). There is no way to identify a unique μ as "the beliefs" of the agent. State-independence is precisely what rules out this indeterminacy.

5 The Main Theorem

Anscombe–Aumann (1963). The preference relation $\succsim$ on H satisfies Axioms 1–5 if and only if there exist a function $v \colon X \to \mathbb{R}$ and a probability distribution $\mu \in \Delta\Omega$ such that

$$U(h) = \sum_{s\in\Omega} \mu(s) \left[\sum_{x\in X} v(x)\, h_s(x) \right] \tag{2}$$

is a utility representation of $\succsim$. Moreover, μ is unique and v is unique up to positive affine transformation, provided there exist $h \succ g$.

The representation (2) has the form of a *subjective* expected utility: the outer sum is an expectation over states under the agent's subjective belief μ, and the inner sum is the vNM expected utility of the roulette lottery h_s under the utility index v. We can write it compactly as:

$$U(h) = \mathbb{E}_\mu \left[\mathbb{E}_{h_s}[v(x)] \right].$$

Expanding the double expectation clarifies the probability accounting:

$$U(h) = \sum_{x \in X} v(x) \underbrace{\left[\sum_{s \in \Omega} \mu(s)\, h_s(x) \right]}_{\Pr(x)},$$

where $\Pr(x) = \sum_s \Pr(s) \Pr(x \mid s, h) = \sum_s \mu(s) h_s(x)$ is the unconditional probability of outcome x under act h and belief μ.

6 Proof of the Representation Theorem

The proof proceeds in four steps. The argument uses the Mixture Space Theorem (equivalently, the vNM theorem) twice: first to obtain a state-dependent additive representation from Axioms 1–4 alone, then to extract subjective probabilities using Axiom 5.

6.1 Step 1 – State-dependent additive representation (Axioms 1–4)

Claim. Axioms 1–4 imply there exist functions $v_1, \dots, v_S \colon X \to \mathbb{R}$ such that

$$U(h) = \sum_{s \in \Omega} \sum_{x \in X} v_s(x)\, h_s(x) \tag{3}$$

represents $\succsim$ on H.

Proof. Since $H = (\Delta X)^\Omega$ with mixing defined pointwise, H is a mixture space. By Axioms 1–4 (completeness, transitivity, independence, Archimedean), the Mixture Space Theorem applies: there exists $F \colon H \to \mathbb{R}$ representing $\succsim$ that is affine, meaning

$$F(\alpha h + (1 - \alpha) g) = \alpha F(h) + (1 - \alpha) F(g). \tag{4}$$

Moreover F is unique up to positive affine transformation.

We now show any such F has the form (3). Fix some reference act $h^* \in H$. For any $h \in H$ and each state s, define $h^s \in H$ to be the act that equals h in state s and equals h^* in all other states. That is, $h^s(s) = h_s$ and $h^s(s') = h^*_{s'}$ for $s' \neq s$.

Observe that

$$\frac{1}{S} h + \frac{S-1}{S} h^* = \sum_{s \in \Omega} \frac{1}{S} h^s$$

since both sides assign, in each state s, the lottery $\frac{1}{S} h_s + \frac{S-1}{S} h^*_s$. Applying (4) and induction:

$$\frac{1}{S} F(h) + \frac{S-1}{S} F(h^*) = \frac{1}{S} \sum_{s \in \Omega} F(h^s). \tag{5}$$

For each s, define $F_s \colon \Delta X \to \mathbb{R}$ by

$$F_s(\pi) = F(h^*_1, \dots, h^*_{s-1}, \pi, h^*_{s+1}, \dots, h^*_S) - \frac{S-1}{S} F(h^*).$$

Then $F_s(h_s) = F(h^s) - \frac{S-1}{S}F(h^*)$. Substituting into (5) and rearranging:

$$F(h) = \sum_{s\in\Omega} F_s(h_s).$$

Since F is affine on H and mixing is pointwise, each F_s is affine on ΔX. Since ΔX is a mixture space of simple probability distributions, any affine function on it has the expected utility form. For each s, define $v_s \colon X \to \mathbb{R}$ by $v_s(x) = F_s(\delta_x)$. The standard induction argument on finite supports gives $F_s(\pi) = \sum_x \pi(x)v_s(x)$. Substituting back establishes (3). □

> *What the axioms are doing.* Axioms 1–4 alone give this state-dependent representation. The key structural move is equation (5): it exploits the pointwise definition of mixing in H to decompose $F(h)$ into a sum of state-specific terms. This decomposition is only possible because an act in H is identified purely by its state-by-state marginals – there is no joint structure across states. This forces additive separability across states before any probability has been identified.

6.2 Step 2 – State-independence forces a common utility index (Axiom 5)

Claim. Under Axiom 5, all non-null states s share the same vNM index up to positive affine transformation: there exist $a_s > 0$ and $b_s \in \mathbb{R}$ such that $v_s(\cdot) = a_s v_{s^*}(\cdot) + b_s$ for any fixed non-null reference state s^*.

Proof. Fix any non-null state s^*. The representation (3) implies that for any $h \in H$ and lotteries $\pi, \rho \in \Delta X$:

$$\sum_x v_s(x)\pi(x) > \sum_x v_s(x)\rho(x) \iff (h_{-s}, \pi) \succ (h_{-s}, \rho).$$

By Axiom 5 (state-independence), for any other non-null state s':

$$(h_{-s}, \pi) \succ (h_{-s}, \rho) \iff (g_{-s'}, \pi) \succ (g_{-s'}, \rho)$$

for all $g \in H$. Therefore, v_s and $v_{s'}$ induce the same ranking of roulette lotteries on ΔX. By the uniqueness part of the Mixture Space Theorem applied to ΔX, if two affine functions represent the same preference, they are positive affine transformations of each other. Hence there exist $a_{s'} > 0$ and $b_{s'} \in \mathbb{R}$ such that

$$v_{s'}(x) = a_{s'} v_{s^*}(x) + b_{s'} \quad \forall x \in X.$$

For null states, v_s is constant (ranking roulette lotteries in a null state never affects overall preferences), so we may set $a_s = 0$ for null states.

> *What Axiom 5 is doing.* Without state-independence, the v_s functions can be completely unrelated across states: one cannot tell whether a high weight $v_s(x)$ reflects a high subjective probability for state s or a genuinely higher utility for x in that state. Axiom 5 rules this out by requiring that the *ranking* of roulette lotteries is the same in every non-null state. This forces all v_s to be positive affine transforms of a single function, resolving the indeterminacy.

6.3 Step 3 – Extract subjective probabilities

Define $U(x) = v_{s^*}(x)$ (normalising to the reference state). Substitute $v_s(x) = a_s U(x) + b_s$ into (3):

$$F(h) = \sum_{s\in\Omega}\sum_{x}(a_s U(x) + b_s)h_s(x) = \sum_{s} b_s + \sum_{s} a_s \left[\sum_{x} U(x)h_s(x)\right].$$

Cancelling the constant $\sum_s b_s$ (which does not affect ordinal rankings) and defining

$$\mu(s) = \frac{a_s}{\sum_{s'\in\Omega} a_{s'}}$$

(which is well-defined since at least one state is non-null, so $\sum_{s'} a_{s'} > 0$), we obtain

$$U(h) \propto \sum_{s\in\Omega} \mu(s) \left[\sum_{x\in X} U(x)h_s(x)\right].$$

By construction $\mu(s) \geq 0$ for all s and $\sum_s \mu(s) = 1$, so μ is a probability distribution over Ω. This establishes the representation (2). □

6.4 Step 4 – Uniqueness

Claim. μ is unique. U is unique up to positive affine transformation.

Proof. Suppose (μ', U') also represents $\succsim$ via (2). Construct the state-dependent functions $V_s(x) = \mu(s)U(x)$ and $V'_s(x) = \mu'(s)U'(x)$. Both give valid state-dependent representations of the form (3). By uniqueness of the Mixture Space Theorem applied to H, there exist $A > 0$ and constants B_s for each s such that $V'_s(\cdot) = AV_s(\cdot) + B_s$, i.e.

$$\mu'(s)U'(x) = A\mu(s)U(x) + B_s \quad \forall x \in X,\ s \in \Omega.$$

Fix a non-null reference state s^0. Since $U(x)$ is non-constant (there exist x, x' with $U(x) \neq U(x')$, as at least two acts are strictly ranked), for each $s \neq s^0$ comparing s and s^0 in this equation gives $B_s = B$ for all s (by the argument of Kreps 1988, p. 110), and $\mu'(s) = A\mu(s) \cdot (U(x') - U(x))/(U'(x') - U'(x))$. Summing over s and using $\sum_s \mu(s) = \sum_s \mu'(s) = 1$ pins down A and forces $\mu' = \mu$. Then $U' = AU + B$, a positive affine transformation. □

> *Why μ is unique but U is not.* The subjective probability μ is pinned down exactly because the roulette lotteries provide an objective ruler: by varying the lotteries assigned to a single state while holding all others fixed, the agent's relative weight on that state is revealed from preferences. No free parameter remains for μ. The utility U retains two degrees of freedom (A and B) – the interval-scale indeterminacy familiar from vNM – because there is no decision-theoretic reason to fix a utility zero or unit.

7 A Worked Example

We trace through the representation for a simple case to make the proof steps concrete.

Primitives. Let $\Omega = \{s_1, s_2, s_3\}$ (three states: say, rain, cloud, sun) and $X = \{x_1, x_2\}$ (two outcomes: umbrella, no umbrella). Roulette lotteries are elements of ΔX, parametrised by a single number $p \in [0, 1]$: $\pi_p = (p, 1-p)$ assigns probability p to x_1.

An act. Consider the act h defined by:

$$h(s_1) = \pi_{0.8}, \quad h(s_2) = \pi_{0.5}, \quad h(s_3) = \pi_{0.1}.$$

In state s_1 (rain), h gives x_1 with probability 0.8; in state s_3 (sun), only with probability 0.1.

Step 1 – State-dependent representation. By Axioms 1–4 and the Mixture Space Theorem, there exist $v_1, v_2, v_3 \colon X \to \mathbb{R}$ such that

$$U(h) = \sum_{s \in \Omega} \big[v_s(x_1)\, h_s(x_1) + v_s(x_2)\, h_s(x_2) \big].$$

Suppose the Mixture Space Theorem delivers $v_1(x_1) = 6$, $v_1(x_2) = 2$, $v_2(x_1) = 3$, $v_2(x_2) = 1$, $v_3(x_1) = 1.5$, $v_3(x_2) = 0.5$.

Then:

$$\begin{aligned} U(h) &= [6(0.8) + 2(0.2)] + [3(0.5) + 1(0.5)] + [1.5(0.1) + 0.5(0.9)] \\ &= [4.8 + 0.4] + [1.5 + 0.5] + [0.15 + 0.45] \\ &= 5.2 + 2.0 + 0.6 = 7.8. \end{aligned}$$

At this stage we have a valid representation, but we cannot say what the agent believes about Ω: any factorisation $v_s(x) = \mu(s) \cdot u(x)$ that reproduces these v_s values is equally valid.

Step 2 – State-independence pins down the common index. Axiom 5 requires that the ranking of roulette lotteries is the same in every non-null state. Note that

$$v_s(x_1) - v_s(x_2) = \begin{cases} 4 & s = s_1 \\ 2 & s = s_2 \\ 1 & s = s_3 \end{cases}$$

These are all positive and all rank x_1 above x_2, consistent with a single underlying index $U(x_1) > U(x_2)$. By uniqueness of the Mixture Space Theorem on ΔX, each v_s must be a positive affine transform of a common U. Fix $U(x_1) = 1$ and $U(x_2) = 0$ (a normalisation). Then:

$$v_s(x_1) = a_s \cdot 1 + b_s, \quad v_s(x_2) = a_s \cdot 0 + b_s = b_s.$$

Reading off: $a_1 = 4$, $b_1 = 2$; $a_2 = 2$, $b_2 = 1$; $a_3 = 1$, $b_3 = 0.5$.

Step 3 – Extract μ. The constants b_s drop out of ordinal comparisons. Define:

$$\mu(s_i) = \frac{a_i}{a_1 + a_2 + a_3} = \frac{a_i}{4 + 2 + 1} = \frac{a_i}{7}.$$

So $\mu(s_1) = 4/7$, $\mu(s_2) = 2/7$, $\mu(s_3) = 1/7$.

Verification. With $U(x_1) = 1$, $U(x_2) = 0$, and these beliefs:

$$\begin{aligned} U(h) &= \mu(s_1)[U(x_1)(0.8) + U(x_2)(0.2)] \\ &\quad + \mu(s_2)[U(x_1)(0.5) + U(x_2)(0.5)] \\ &\quad + \mu(s_3)[U(x_1)(0.1) + U(x_2)(0.9)] \\ &= \tfrac{4}{7}(0.8) + \tfrac{2}{7}(0.5) + \tfrac{1}{7}(0.1) \\ &= \tfrac{3.2+1.0+0.1}{7} = \tfrac{4.3}{7} \approx 0.614. \end{aligned}$$

This is an affine transformation of the 7.8 computed above (subtracting $\sum_s b_s = 2+1+0.5 = 3.5$ and dividing by 7 gives $(7.8-3.5)/7 = 4.3/7$), confirming the representations are ordinally equivalent.

What the example shows. The state-dependent functions v_s in Step~1 confound beliefs and tastes: one cannot tell from $v_1(x_1) = 6$ whether the agent values x_1 highly or simply thinks state s_1 is likely. State-independence (Step~2) disentangles the two by forcing all v_s to share a common shape. The a_s weights then carry all the belief content, and normalising them to sum to one delivers the unique subjective probability μ.

8 What A–A Adds to vNM

	vNM	**Anscombe–Aumann**
Domain	Lotteries ΔX	Acts $H = (\Delta X)^\Omega$
Probabilities	Objective, given	Subjective, derived from $\succsim$
What $\succsim$ identifies	v (up to affine)	μ (uniquely) + v (up to affine)
Key extra axiom	–	State-independence
Representation	$\sum_x v(x)\pi(x)$	$\sum_s \mu(s) \sum_x v(x) h_s(x)$

The A–A framework strictly extends vNM. The constant acts $H_c \cong \Delta X$ form a subset of H: restricting A–A preferences to H_c recovers the vNM setting exactly. The extra structure – the state space Ω and the state-independence axiom – is what allows beliefs to be separated from tastes and identified uniquely.

9 What Each Axiom Contributes

Axiom	Role in the proof
Completeness, Transitivity	Guarantee a complete, transitive ranking of acts; necessary for any representation.
Independence	Applied to H: forces affine (linear-in-probabilities) structure on the representation. On constant acts, recovers vNM independence. On non-constant acts, forces additivity across states.
Archimedean	Rules out lexicographic preferences over acts; ensures continuous mixture-space structure, allowing the Mixture Space Theorem to apply.
State-Independence	The key additional axiom. Forces the vNM index v to be the same function in every non-null state; separates utility from beliefs; enables unique identification of μ. Without it, only a state-dependent representation (Theorem 1) is possible and μ cannot be uniquely recovered.

10 References

- Anscombe, F. J. and Aumann, R. J. (1963). "A definition of subjective probability." *Annals of Mathematical Statistics* 34(1): 199–205.
- Gilboa, I. (2009). *Theory of Decision Under Uncertainty.* Cambridge University Press.
- Kreps, D. (1988). *Notes on the Theory of Choice.* Westview Press.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory.* Oxford University Press.
- Savage, L. J. (1954). *The Foundations of Statistics.* Wiley.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior.* Princeton University Press.