

Choice under risk

Robert J. Carroll

Last revised: 6 May 2026

1 Motivation

The previous handout developed the theory of choice over deterministic alternatives: a voter selects a candidate, an agenda setter selects a motion, and the consistency of those selections is captured by WARP and SARP. Most of the choices political-economy models actually represent are not deterministic. A candidate chooses a campaign strategy whose payoff depends on turnout. A legislator votes for a reform whose effects depend on legislative state-of-the-world that hasn't yet revealed itself. A bureaucrat allocates effort across portfolios with uncertain rewards. Each of these is a choice not over an outcome, but over a *gamble* — a probability distribution over outcomes.

This handout develops the formal apparatus for choice over gambles when the probabilities are *given*. The framework is due to Neumann and Morgenstern (1944): the alternative space is replaced by the set of probability distributions over the outcome space, and the preference relation is now over those distributions. Under four axioms — completeness, transitivity, continuity, and an axiom called independence that does the load-bearing work — there is a representation theorem that recovers a real-valued utility function on outcomes such that preferences over lotteries are determined by expected utility. The theorem is one of the cleanest in the formal social sciences. It is also one whose central axiom, independence, fails empirically in a robust way — the failure is the subject of the Allais paradox in §5 and of prospect theory in §6.

The honest framing for this handout is that expected-utility theory is right where it is right and wrong where it is wrong, and the precise location of each is what we want. The theorem is exactly what is claimed: under those four axioms, expected-utility maximization is forced. Empirical preferences sometimes satisfy the four axioms and sometimes don't. When they don't, the failure is structural and worth understanding on its own terms. We get to both halves — the theorem, and the structural failures it permits — before the handout closes.

2 Lotteries

When a candidate weighs a campaign strategy, what she is choosing is not an outcome — she does not yet know the outcome. She is choosing a probability distribution over outcomes. The same is true of the legislator weighing a reform with uncertain effects, the regulator weighing a policy with stochastic compliance, and the bureaucrat allocating effort across portfolios with uncertain returns. Most of the choices political-economy models actually represent have this structure. Before we can talk about preferences over those probability distributions, we need a name for the object the agent is choosing and a small amount of formal scaffolding to keep its structure straight. The object is the *lottery*, and the scaffolding is one section long; we then move on.

We restrict to finite outcome spaces X throughout the main development; the general measurable case (probabilities on a measurable space, expectations as Lebesgue integrals) uses the same machinery

and is covered by the probability cluster, but the finite case is enough to display every substantive idea.¹

Definition 1. A *lottery* (or *simple lottery*) over a finite outcome space X is a function $p : X \rightarrow [0, 1]$ with $\sum_{x \in X} p(x) = 1$. The set of all lotteries over X is

$$\Delta(X) = \left\{ p : X \rightarrow [0, 1] : \sum_{x \in X} p(x) = 1 \right\},$$

which is a simplex in \mathbb{R}^X of dimension $|X| - 1$.

For $|X| = 3$ this is a triangle in \mathbb{R}^3 (the Marschak–Machina triangle, which we will use in a footnote later); for $|X| = 2$ it is a line segment.

A *compound lottery* over X is a probability distribution over lotteries: a probability α of receiving lottery p and probability $1 - \alpha$ of receiving lottery q , where $p, q \in \Delta(X)$. The natural compound is denoted $\alpha p + (1 - \alpha)q$. Its *reduced* (or *simple*) form is the lottery in $\Delta(X)$ obtained by simply weighting:

$$[\alpha p + (1 - \alpha)q](x) = \alpha p(x) + (1 - \alpha)q(x).$$

Assumption 2 (Reduction of compound lotteries). The agent’s preference is over reduced lotteries, so the compound $\alpha p + (1 - \alpha)q$ is treated as identical to its reduced form. Equivalently: the agent does not have additional preferences over the timing or layered structure of the gamble.

The reduction assumption is the technical move that lets us treat $\Delta(X)$ as the universe of choice objects. It is innocent for most applied purposes; the cases where it fails (preferences over timing of resolution, the Kreps–Porteus framework) are real but beyond this handout.

Example 3 (Campaign strategy as a lottery). A candidate considers two campaign strategies. The *safe* strategy is a turnout-mobilization push: with probability 0.7 the candidate gains 5 percentage points, with probability 0.3 she gains 0. The *moonshot* strategy is a high-profile persuasion campaign aimed at swing voters: with probability 0.2 she gains 20 points, with probability 0.5 she gains 5, with probability 0.3 she loses 5. The outcome space is $X = \{-5, 0, 5, 20\}$ (the gain in percentage points), and the two strategies correspond to two distinct points in $\Delta(X)$.

We will return to this example to display each piece of the framework as it is built. The vNM theorem, in its applied form, is the assertion that the candidate’s choice between the two strategies is determined — given her preferences satisfy the four axioms — by computing $\mathbb{E}[u]$ under each lottery and picking the larger.

¹Generalizing to infinite outcome spaces is a question of measure-theoretic bookkeeping rather than conceptual depth. With X a measurable space, $\Delta(X)$ is the space of probability measures on X ; expected utility is $\mathbb{E}_p[u] = \int u dp$, and the vNM axioms below have natural reformulations (continuity becomes a topological condition on $\Delta(X)$ in the weak-* topology). The probability spaces and measures handout, the integration handout, and the random variables and expectations handout supply all the machinery needed. The vNM theorem extends, modulo regularity conditions on u (boundedness or domain restrictions), to the infinite case; see Kreps (1988) or Mas-Colell, Whinston, and Green (1995, §6.B).

3 The von Neumann–Morgenstern axioms

What conditions on the candidate’s preferences over campaign-strategy lotteries would let us write her choice as a calculation: assign each outcome a real number, take the expectation under each lottery, pick the larger? The question is not idle. Most of applied political economy already proceeds as if the agent’s preferences could be encoded that way — expected utility shows up unannounced in turnout models, in deterrence games, in legislative bargaining, in models of regulatory enforcement. The von Neumann–Morgenstern theorem is the disciplined statement of when that move is licensed, and the four axioms below are exactly the conditions it requires.

A preference relation \succsim on $\Delta(X)$ is a complete preorder, with strict part \succ and indifference part \sim inherited as in the previous handouts. Two of the four vNM axioms are the order-theoretic axioms we already know — completeness and transitivity. Two are new: continuity (which rules out lexicographic-style preferences over lotteries) and independence (which is the load-bearing axiom and the one with empirical bite). The rest of this handout is mostly the story of what independence does, and where it fails.

Axiom 4 (Completeness). For every $p, q \in \Delta(X)$, either $p \succsim q$ or $q \succsim p$.

Axiom 5 (Transitivity). For every $p, q, r \in \Delta(X)$: $p \succsim q$ and $q \succsim r$ imply $p \succsim r$.

These two are the same axioms that make \succsim a complete preorder, restated in the lottery setting. Their content is identical to the deterministic case from order theory.

Axiom 6 (Continuity (Archimedean)). For every $p, q, r \in \Delta(X)$ with $p \succ q \succ r$, there exist $\alpha, \beta \in (0, 1)$ such that

$$\alpha p + (1 - \alpha)r \succ q \quad \text{and} \quad q \succ \beta p + (1 - \beta)r.$$

The reading: between any strictly-preferred pair $p \succ r$, the probability mixtures of p and r trace out a continuous curve from r (when the weight on p is zero) to p (when the weight on p is one), and any intermediate $q \in (r, p)$ in the strict-preference order can be straddled by mixtures from both sides. Equivalently: there are no infinitely-preferred outcomes that swamp arbitrarily small probabilities of other outcomes (no lexicographic preferences over lotteries). The condition rules out, e.g., a preference that puts “life with probability 1” strictly above “life with probability 0.99 and \$100 with probability 0.01” and refuses to compare them via a finite mixture.

Axiom 7 (Independence). For every $p, q, r \in \Delta(X)$ and every $\alpha \in (0, 1]$:

$$p \succsim q \iff \alpha p + (1 - \alpha)r \succsim \alpha q + (1 - \alpha)r.$$

Independence is the load-bearing axiom of the vNM framework, and the one that empirically fails. The reading: if p is preferred to q , then mixing both with the *same* third lottery r , in the same proportions, preserves the preference. The third lottery’s contribution is a “common consequence” that is irrelevant to the comparison between $\alpha p + (1 - \alpha)r$ and $\alpha q + (1 - \alpha)r$, because in both cases you face r with the same probability and the original lottery with the same probability. The axiom asks the agent to disregard the common consequence in evaluating the comparison.

If independence holds for all r and α , the indifference curves on $\Delta(X)$ are forced to be parallel hyperplanes. (This is the geometric content of independence.) That parallel-hyperplane structure is exactly what produces the linearity-in-probabilities of the expected-utility representation in the next section, and it is what fails in the empirical Allais pattern of §5.

4 The expected-utility theorem

The vNM axioms set up a question with a clean answer. Suppose the candidate's preference over campaign-strategy lotteries satisfies all four. What does her preference look like? The theorem below gives exactly the calculation we hoped for in the previous section: a single utility function on outcomes (the *vNM utility*), with her preference over any two lotteries determined by computing expected utility under each and picking the larger. The proof is short and the structure is illuminating: each of the four axioms does a specific piece of work, and seeing which axiom does what is the cleanest way to see why independence is the one that matters.

Theorem 8 (von Neumann–Morgenstern). *Let X be finite. A complete preorder \succsim on $\Delta(X)$ satisfies continuity and independence iff there exists a function $u : X \rightarrow \mathbb{R}$ such that, for every $p, q \in \Delta(X)$,*

$$p \succsim q \iff \mathbb{E}_p[u] \geq \mathbb{E}_q[u],$$

where $\mathbb{E}_p[u] = \sum_{x \in X} p(x)u(x)$. Moreover, u is unique up to positive affine transformation: if u' also represents \succsim , then $u' = au + b$ for some $a > 0$ and $b \in \mathbb{R}$.

The function u is called the *vNM utility* on outcomes. The proof is the cleanest illustration of how the four axioms wire together. We give it in three steps.

Proof of the representation, finite X . **Step 1: extract a best and worst outcome.** Since \succsim is complete and transitive on $\Delta(X)$ and X is finite, there exist $\bar{x}, \underline{x} \in X$ that are \succsim -best and -worst among the degenerate lotteries (the lotteries that put probability 1 on a single outcome). Without loss of generality assume $\bar{x} \succ \underline{x}$ (else \succsim is trivial and any constant u represents it).

Step 2: extract an indifference probability for each outcome. For each $x \in X$, the degenerate lottery $\delta_x \in \Delta(X)$ satisfies $\bar{x} \succ \delta_x \succ \underline{x}$. By continuity (Axiom 6, applied with $p = \delta_{\bar{x}}$, $q = \delta_x$, $r = \delta_{\underline{x}}$), there exists a unique $\alpha_x \in [0, 1]$ such that

$$\delta_x \sim \alpha_x \delta_{\bar{x}} + (1 - \alpha_x) \delta_{\underline{x}}.$$

(Uniqueness: if two values $\alpha_x \neq \alpha'_x$ both produced indifference, independence applied to the difference would generate a strict preference, contradicting indifference. Existence: continuity provides α_x via an intermediate-value argument on the line of mixtures from $\delta_{\underline{x}}$ to $\delta_{\bar{x}}$.) Define $u(x) := \alpha_x$.

Step 3: extend u linearly to lotteries via independence. For any $p \in \Delta(X)$, the reduction-of-compound-lotteries assumption and repeated application of independence yield

$$p \sim \left(\sum_{x \in X} p(x) \alpha_x \right) \delta_{\bar{x}} + \left(1 - \sum_{x \in X} p(x) \alpha_x \right) \delta_{\underline{x}}.$$

This is the linearity step: the indifference probability of p between $\delta_{\bar{x}}$ and $\delta_{\underline{x}}$ is the p -weighted average of the per-outcome indifference probabilities α_x . Since \succsim on the line of mixtures from $\delta_{\bar{x}}$ to $\delta_{\underline{x}}$ is determined by the coefficient on $\delta_{\bar{x}}$ (a higher coefficient means a strictly preferred lottery, by independence and $\bar{x} \succ \underline{x}$), we have

$$p \succsim q \iff \sum_x p(x) \alpha_x \geq \sum_x q(x) \alpha_x \iff \mathbb{E}_p[u] \geq \mathbb{E}_q[u].$$

Affine uniqueness. If u' also represents \succsim , the function $\alpha \mapsto \mathbb{E}_p[u']$ on the line $p = \alpha \delta_{\bar{x}} + (1 - \alpha) \delta_{\underline{x}}$ is monotone (because \succsim -monotone on the line) and linear (by reduction + independence). A monotone

linear function on $[0, 1]$ is determined up to positive affine transformation. Hence $u' = au + b$ for some $a > 0$, $b \in \mathbb{R}$.

Necessity. The expected-utility representation \succsim is automatically complete, transitive, continuous, and independent: the four axioms are properties of $\mathbb{E}[u]$ as a real-valued functional on $\Delta(X)$ that is linear in p . \square

The theorem is doing two things at once. It is a representation result: every preference satisfying the four axioms is rationalized by an expectation of a utility function. And it is an axiomatization: the four axioms are exactly the content of expected-utility maximization, with no slack. The independence axiom is the one that produces the *linearity* of the representation in probabilities, and removing or weakening it breaks linearity in characteristic ways.²

5 The Allais paradox

The vNM theorem is a beautiful piece of mathematics, and its central axiom is the one that empirical preferences violate. The cleanest demonstration is over sixty years old now. Maurice Allais, in 1953, presented experimental subjects with two pairs of lotteries on a three-outcome space and elicited a modal pattern of preferences that no expected-utility maximizer can hold simultaneously (Allais, 1953). The pattern is robust. It has been replicated everywhere the framework has been tried — across populations, currencies, decades, and political-economy settings where the lotteries are over policy outcomes or candidate-victory probabilities. We work through the original example.

The outcome space is $X = \{\$0, \$1\text{M}, \$5\text{M}\}$ (in nominal millions of US dollars; the magnitudes are not the substantive content, but they make the lotteries vivid). The four lotteries:

	\$0	\$1M	\$5M
L_1	0	1	0
L_2	0.01	0.89	0.10
L_3	0.89	0.11	0
L_4	0.90	0	0.10

The modal pattern of preferences in repeated experiments is $L_1 \succ L_2$ (the certain million is preferred to a small chance of getting nothing in exchange for a smaller chance of five million) and $L_4 \succ L_3$ (when both lotteries already have a high chance of zero, the higher-EV option is appealing). The two pair-comparisons together violate independence.

²It is worth pausing on what independence is geometrically doing. Indifference curves on the simplex $\Delta(X)$, under the four vNM axioms, are parallel hyperplanes — their gradient is the vector $(u(x))_{x \in X}$. “Parallel” is the geometric content of independence: if you mix two indifferent lotteries with the *same* third lottery in the *same* proportions, the resulting two lotteries are still indifferent. The Marschak–Machina triangle is the simplex $\Delta(X)$ for $|X| = 3$, drawn as a triangle in the plane, and the parallel-hyperplane structure is visible immediately: indifference curves are parallel line segments. Empirically, the Allais pattern in §5 produces indifference curves that *fan out* from a corner of the triangle rather than running parallel — the curvature is exactly the failure of independence. Machina (1982) works out the consequences of relaxing independence while keeping the rest of the framework: “expected utility analysis without the independence axiom” becomes a smooth-utility theory in which preferences are locally linear (so most of the comparative-statics machinery survives) but globally nonlinear (so Allais is consistent). Machina’s framework is the principled non-EU response that is distinct from prospect theory: it weakens the axiom system at the structural level, while prospect theory rebuilds the representation from psychological primitives.

Proposition 9. *No vNM preference satisfies $L_1 \succ L_2$ and $L_4 \succ L_3$.*

Proof. Suppose, for contradiction, that $u : X \rightarrow \mathbb{R}$ represents \succsim in the vNM sense. Writing $u_0 = u(\$0)$, $u_1 = u(\$1M)$, $u_5 = u(\$5M)$:

- $L_1 \succ L_2$ says $u_1 > 0.01u_0 + 0.89u_1 + 0.10u_5$, which simplifies to $0.11u_1 > 0.01u_0 + 0.10u_5$.
- $L_4 \succ L_3$ says $0.90u_0 + 0.10u_5 > 0.89u_0 + 0.11u_1$, which simplifies to $0.01u_0 + 0.10u_5 > 0.11u_1$.

The two inequalities are contradictory. □

The structural reading: L_2 is built from L_1 by replacing 0.11 probability of \$1M with a (0.01, 0.10) split between \$0 and \$5M. L_4 is built from L_3 by the same replacement: 0.11 probability of \$1M replaced with the same (0.01, 0.10) split. The two replacements are identical, so independence requires the two pair-preferences to go in the same direction. They don't.

The Allais pattern is robust. It is reproduced across populations, currencies, and political-economy settings (lottery comparisons over policy outcomes, candidate-victory probabilities, regulatory enforcement gambles). The conclusion is that empirical preferences over lotteries are not, in general, vNM. The next section turns to the most prominent positive theory of what they are instead.

6 Prospect theory

The Allais pattern is robust enough that some positive theory of what's going on is wanted — a successor to expected utility that gets the empirical pattern right and keeps the theory's basic shape (a real-valued functional that an agent maximizes). The most influential such successor is *prospect theory*, developed by Kahneman and Tversky (1979) and refined repeatedly since. Prospect theory is unapologetically psychological. It takes seriously the experimental finding that people respond to gains and losses asymmetrically, that they overweight small probabilities and underweight large ones, and that the reference point against which gains and losses are measured matters substantively to the agent's choices. The result is a full functional alternative to expected utility, with three components: a *value function* v defined relative to a reference point; a *probability weighting function* w on the per-outcome probabilities; and a *loss-aversion* kink at the reference point. The full representation is

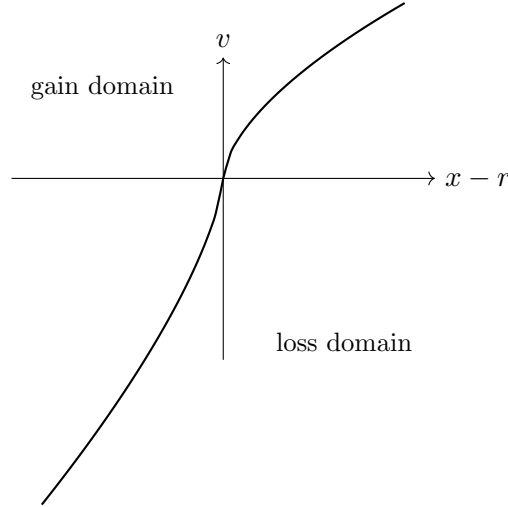
$$V(p) = \sum_{x \in X} w(p(x)) \cdot v(x - r),$$

where r is the reference point.

The pieces:

- *Value function* v . Defined on differences from the reference point. Concave for gains ($x > r$, so $v(x - r)$ exhibits diminishing marginal returns), convex for losses ($x < r$, so additional losses have diminishing marginal pain), with a kink at 0. The standard parametric form is $v(z) = z^\alpha$ for $z \geq 0$ and $v(z) = -\lambda(-z)^\beta$ for $z < 0$, with $\alpha, \beta \in (0, 1)$ and $\lambda > 1$ (loss aversion). The S-shaped value function:

Reading the figure: v is concave on the right of zero, convex on the left, and steeper on the left (the loss-aversion asymmetry).



- *Probability weighting w .* A function $w : [0, 1] \rightarrow [0, 1]$ with $w(0) = 0$, $w(1) = 1$, that overweights small probabilities and underweights large ones. The standard parametric form is $w(p) = p^\gamma / (p^\gamma + (1 - p)^\gamma)^{1/\gamma}$ for $\gamma \approx 0.61$, producing the inverse-S shape.
- *Reference point r .* The point of comparison. In most applications this is the status quo: the agent's current wealth, the incumbent legislator's current rule, the existing policy. Formally exogenous in the basic theory; endogenizing r is an active research direction.

The PT functional explains the Allais pattern: the overweighting of small probabilities makes the 0.01-probability of \$0 in L_2 feel like a larger threat than its expected value warrants, pushing $L_1 \succ L_2$; the same overweighting also makes the 0.10-probability of \$5M in L_4 feel especially appealing relative to a base of zero, pushing $L_4 \succ L_3$. The independence axiom is rejected by the structure of w (a nonlinear weighting on probabilities is exactly the failure of linearity-in-probabilities that independence enforces).

PT is therefore not in the EU class. The functional V is not the expectation of any utility function: w being nonlinear means $w(p(x))$ does not assemble into a probability measure, and the representation is no longer an integral.³

³The endogeneity question for the reference point r is the most substantive lacuna in PT and the place where the political-economy applications get interesting. The standard reading in applied work is that r is the *status quo*: the existing policy, the incumbent's record, the agent's current portfolio. Quattrone and Tversky (1988) is the canonical political-science demonstration, showing that voters frame policy outcomes relative to the status quo and exhibit risk aversion in the gain domain (when the proposed change is a gain) but risk seeking in the loss domain (when the proposed change is framed as preventing a loss). The political-economy implication is that reform proposals are systematically harder to pass than expected-utility theory would predict, and the asymmetry is empirically detectable across legislatures and electorates. Endogenizing r has been pursued in several directions: Kőszegi and Rabin (2006) build expectations-based reference points (the agent's reference is what she expected to receive, calibrated by the equilibrium of a dynamic process), and Tversky and Kahneman (1992) build cumulative prospect theory (which weights cumulative probabilities, restoring stochastic-dominance respect that the basic PT functional can violate). Both extensions are active areas; this handout flags them but does not pursue them.

7 Rank-dependent expected utility

Prospect theory is one option for a successor framework, but it has a known problem: the per-outcome probability weighting $w(p(x))$ can be aggressive enough that the resulting functional violates first-order stochastic dominance — the agent strictly prefers a lottery that is dominated by another, outcome by outcome and probability by probability. That is hard to defend as an account of rational choice. There is a cleaner weakening of independence that preserves stochastic dominance, and it is worth knowing because it is the structurally honest middle ground between expected utility and prospect theory. *Rank-dependent expected utility*, due to Quiggin (1982), applies the probability weighting not to each outcome’s probability separately but to the *cumulative* probabilities of outcomes ranked from worst to best.

Concretely: order $X = \{x_{(1)}, \dots, x_{(n)}\}$ from worst to best by \succsim , and let $F_p(k) = \sum_{i \leq k} p(x_{(i)})$ be the cumulative probability of receiving outcome $x_{(k)}$ or worse. The rank-dependent functional is

$$\text{RDEU}(p) = \sum_{k=1}^n [w(1 - F_p(k-1)) - w(1 - F_p(k))] \cdot u(x_{(k)}),$$

where w is a probability weighting function applied to the *decumulative* probability $1 - F_p$. The weight on outcome $x_{(k)}$ is the difference of two weighted decumulative probabilities, ensuring the weights sum to 1 and that better outcomes have weights determined by the cumulative shape of w .

The honest story:

- RDEU agrees with EU when w is the identity ($w(p) = p$): the difference $w(1 - F_p(k-1)) - w(1 - F_p(k)) = p(x_{(k)})$, giving back the standard expected utility.
- RDEU agrees with EU on *two-outcome* lotteries even when w is nonlinear: with only two outcomes $x_{(1)} \prec x_{(2)}$, the weights collapse to $w(p(x_{(2)}))$ on the better outcome and $1 - w(p(x_{(2)}))$ on the worse. Linearity-in-probabilities is preserved when there are only two outcomes; non-linearity manifests starting at three outcomes.
- RDEU *respects first-order stochastic dominance*: if p first-order stochastically dominates q , then $\text{RDEU}(p) \geq \text{RDEU}(q)$ for any monotone w . This is the property the basic PT functional fails when w is sufficiently curved.
- RDEU still violates independence (the cumulative weighting is nonlinear in p in general), so it is not in the EU class. But it is a principled weakening that keeps the dominance-respecting structure that EU has and PT doesn’t.

Tversky and Kahneman (1992) build *cumulative* prospect theory, a hybrid that puts RDEU’s cumulative weighting inside the PT framework: gain-domain outcomes are weighted by one cumulative-weighting function, loss-domain outcomes by another, and the value function is still S-shaped with a loss-aversion kink. CPT preserves stochastic dominance and is the contemporary applied version of prospect theory.

The map of the EU-and-its-weakenings landscape, with each model named by what it does to the independence axiom:

The columns name the two structural conditions a representation can have. The first column is the diagnostic that distinguishes EU from non-EU; the second is the diagnostic that distinguishes principled weakenings (RDEU, CPT, Machina) from less principled ones (basic PT, where the weighting can dominate the value function and violate dominance for sufficiently curved w).

Model	Independence	Stochastic dominance
Expected utility (vNM)	holds	respected
Rank-dependent EU	weakened	respected
Cumulative prospect theory	weakened	respected
Prospect theory (1979)	weakened	can be violated
Machina's smooth utility	local only	respected

8 What's next

The next handout drops the assumption that probabilities are given. The framework changes: alternatives are *acts*, mappings from states of the world to outcomes (Savage) or to lotteries (Anscombe–Aumann), and the agent's preference is over acts. The Anscombe–Aumann framework explicitly builds on the lottery machinery from this handout, and the modern axiomatic treatments of ambiguity — Gilboa–Schmeidler maxmin EU, Schmeidler's Choquet EU — live in that framework. The Ellsberg paradox, the canonical empirical failure of subjective probability, is the analogue of Allais one level up.

For the vNM theorem and its applied use, Kreps (1988) and Mas-Colell, Whinston, and Green (1995, §6) are the standard graduate references. Gilboa (2009) is more recent and treats EU and ambiguity together. Kahneman and Tversky (1979) and Tversky and Kahneman (1992) are the source papers for prospect theory; Quiggin (1982) for rank-dependent EU; Machina (1982) for the smooth-utility weakening. Quattrone and Tversky (1988) is the canonical political-science demonstration of PT in voter behavior. Austen-Smith and Banks (1999) treats the EU framework throughout with PE applications.

9 Exercises

Exercise 10. Verify the affine-uniqueness claim of the vNM theorem on a small example. Let $X = \{x_1, x_2, x_3\}$ and suppose $u(x_1) = 0$, $u(x_2) = 1$, $u(x_3) = 4$ represents \succsim . (a) Show that $u'(x) = 2u(x) + 5$ also represents \succsim . (b) Show that $u''(x) = u(x)^2$ generally does not (find two lotteries on which $\mathbb{E}[u]$ and $\mathbb{E}[u'']$ disagree).

Exercise 11. Compare the two campaign strategies of Example 3 under expected-utility maximization with two different vNM utility functions: (a) risk-neutral, $u(x) = x$, and (b) risk-averse, $u(x) = \log(x + 10)$ (the offset ensures u is defined on the loss outcome $x = -5$). Which strategy does each utility prefer? What does the comparison say about the relationship between the curvature of u and the candidate's tolerance for high-variance strategies?

Exercise 12. Prove that under the vNM axioms, the indifference curves on $\Delta(X)$ are parallel hyperplanes. (Hint: the gradient of $\mathbb{E}_p[u]$ with respect to the probability vector p is the vector $(u(x))_{x \in X}$, which does not depend on p .)

Exercise 13. *Allais's paradox, written out.* Verify Proposition 9 by direct computation: assume a vNM utility $u : \{\$0, \$1M, \$5M\} \rightarrow \mathbb{R}$ exists, suppose both $L_1 \succ L_2$ and $L_4 \succ L_3$ hold, and derive the contradiction.

Exercise 14. *Reform proposal as a lottery.* A legislator faces a vote on a reform whose policy effects are uncertain. The status quo (incumbent policy) gives outcome $\$x_0$ with certainty, where

x_0 is the legislator’s reference point. The reform gives outcome $\$x_+ > x_0$ with probability α and $\$x_- < x_0$ with probability $1 - \alpha$. Under the basic PT value function $v(z) = z^{0.88}$ for $z \geq 0$ and $v(z) = -2.25(-z)^{0.88}$ for $z < 0$ (the Tversky–Kahneman 1992 calibration), and unweighted probabilities (w identity), find the threshold value of α above which the legislator strictly prefers reform to status quo. Compare to the threshold under risk-neutral expected-value maximization. Discuss in one or two sentences how loss aversion makes reform politically harder than expected-value reasoning would predict.

Exercise 15. *Voter risk preferences and Allais.* Two policy proposals are framed as lotteries over outcomes (an ambitious-but-uncertain plan vs. a modest-but-certain plan). Suppose voter preferences exhibit the same Allais pattern as the lab experiments: when the modest plan is itself a lottery (e.g., subject to legislative haggling), voters prefer the ambitious plan; when the modest plan is certain (e.g., a binding pre-commitment), voters prefer it. Set up two pairs of lotteries that exhibit this Allais-style pattern, and verify that no vNM preference rationalizes both of the modal pair preferences.

Exercise 16. Show that rank-dependent EU on a two-outcome lottery coincides with EU. (Hint: with $X = \{x_{(1)}, x_{(2)}\}$ and $x_{(1)} \prec x_{(2)}$, the cumulative probabilities are $F_p(1) = p(x_{(1)})$, $F_p(2) = 1$. Compute $\text{RDEU}(p)$ and show it equals $w(p(x_{(2)}))u(x_{(2)}) + (1 - w(p(x_{(2)})))u(x_{(1)})$, which is a vNM expectation under the relabeled probability $\tilde{p}(x_{(2)}) = w(p(x_{(2)}))$.)

Exercise 17. Give a three-outcome lottery on which RDEU and basic prospect theory disagree about a stochastic-dominance comparison. (Hint: pick a w with sufficient curvature that the per-outcome PT weights $w(p(x))$ assign more weight to a worst outcome than its probability warrants. The cumulative-weighting structure of RDEU avoids this; the per-outcome structure of basic PT does not.)

Exercise 18. Suppose a candidate’s vNM utility over campaign outcomes (in vote-share gain) is $u(x) = -e^{-\rho x}$ for $\rho > 0$ (constant absolute risk aversion). (a) Show that the candidate’s preference between any two campaign-strategy lotteries is unchanged when the same constant gain c is added to every outcome. (b) Show that the candidate’s preference depends on ρ but not on ρ ’s sign in a way that is monotone: a higher ρ corresponds to a more risk-averse candidate. (c) Discuss in one sentence why CARA is a particularly clean (and unrealistic) functional form for political-economy modeling.

Exercise 19. A small TikZ-free verification. The Marschak–Machina triangle is the simplex $\Delta(X)$ for $|X| = 3$, drawn in the plane with the three corners labeled by the three degenerate lotteries. Place the four Allais lotteries L_1, L_2, L_3, L_4 in the triangle (i.e., write down their barycentric coordinates with respect to the corners $\delta_{\$0}, \delta_{\$1M}, \delta_{\$5M}$). Then identify the two pairs that share a common direction of difference — the structural reason the modal Allais pattern violates independence.

References

- Allais, Maurice (1953). “Le Comportement de l’Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l’École Américaine”. In: *Econometrica* 21.4, pp. 503–546.
- Austen-Smith, David and Jeffrey S. Banks (1999). *Positive Political Theory I: Collective Preference*. Ann Arbor: University of Michigan Press.
- Gilboa, Itzhak (2009). *Theory of Decision under Uncertainty*. Cambridge: Cambridge University Press.

- Kahneman, Daniel and Amos Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk". In: *Econometrica* 47.2, pp. 263–291.
- Kőszegi, Botond and Matthew Rabin (2006). "A Model of Reference-Dependent Preferences". In: *Quarterly Journal of Economics* 121.4, pp. 1133–1165.
- Kreps, David M. (1988). *Notes on the Theory of Choice*. Boulder, CO: Westview Press.
- Machina, Mark J. (1982). "“Expected Utility” Analysis without the Independence Axiom". In: *Econometrica* 50.2, pp. 277–323.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995). *Microeconomic Theory*. New York: Oxford University Press.
- Neumann, John von and Oskar Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- Quattrone, George A. and Amos Tversky (1988). "Contrasting Rational and Psychological Analyses of Political Choice". In: *American Political Science Review* 82.3, pp. 719–736.
- Quiggin, John (1982). "A Theory of Anticipated Utility". In: *Journal of Economic Behavior & Organization* 3.4, pp. 323–343.
- Tversky, Amos and Daniel Kahneman (1992). "Advances in Prospect Theory: Cumulative Representation of Uncertainty". In: *Journal of Risk and Uncertainty* 5.4, pp. 297–323.