Choice under ambiguity

Robert J. Carroll

Last revised: 6 May 2026

---

1 Motivation

The previous handout assumed the agent knew the probabilities. A great many political-economy decisions don't have known probabilities, and don't have probabilities the agent could plausibly elicit even by introspection. A regulator setting climate policy faces several competing scientific models of how warming responds to emissions, with no agreed-upon prior over which model is correct. A war planner contemplating a novel adversary — new technology, unfamiliar alliance configuration, no recent track record — has no frequency-based prior on the adversary's resolve type. A constitutional designer choosing institutional rules cannot put a probability on which kinds of political coalitions will form under those rules over the next century, because the rules themselves help determine what coalitions will form. In each of these cases, the probabilities themselves are part of what the agent is uncertain about, and the framework of the previous handout doesn't apply.

The decision-theoretic literature has three positions on what to do, in increasing order of departure from the Bayesian programme. The first is Savage's: probabilities can always be elicited from the agent's preferences over bets, and a sufficiently rational agent has a unique subjective probability whether or not she can articulate it. The second is Knight's, reopened by Ellsberg: there is a substantive gap between subjective probability and the agent's actual epistemic state, and empirical preferences detect the gap. The third, formalized by Gilboa and Schmeidler in the late 1980s, is the modern multi-prior view: the agent has a *set* of priors and a decision rule (typically the maxmin) for choosing among acts. This handout takes each of the three in turn, with the multi-prior framework as the centerpiece because it is the framework in which most contemporary applied work on ambiguity actually proceeds.

The handout's order is built around a single pedagogical move. Savage's framework is the cleanest statement of the Bayesian position but technically heavy. The Anscombe–Aumann framework, which builds explicitly on the lottery machinery of the previous handout, gives the same conclusion (a unique probability, a utility unique up to positive affine transformation) with much lighter axiomatic lifting. The modern multi-prior models all live in the Anscombe–Aumann framework, so spending a section on it before Ellsberg is what makes the rest of the handout legible.

2 Savage's subjective expected utility

The Bayesian programme says: probabilities live in the agent, not in the world, and the agent's choices among bets reveal them. Savage's 1954 *Foundations of Statistics* is the most ambitious axiomatic vindication of this view, deriving a unique subjective probability and a utility on outcomes from preferences over a different kind of object: *acts*, which are mappings from states-of-the-world to outcomes.

Definition 1. A *Savage act* is a function $f : S \to X$, where S is a set of *states of the world* (with a σ-algebra Σ) and X is a set of *outcomes*. The agent's preferences are over acts: $f \succsim g$ means the agent weakly prefers act f to act g.

Acts encode the data of a decision under uncertainty. The state is whatever uncertainty is relevant to the agent's choice (whether the climate model is high-sensitivity or low-sensitivity, whether the adversary's resolve is high or low, whether the proposed institutional rule produces stable or unstable coalitions). The outcome is what the agent ends up with, given the state and her chosen act.

Savage's framework imposes seven axioms on $\succsim$, conventionally numbered P1–P7. P1 is completeness and transitivity. P2 (the *sure-thing principle*) is the act-level analogue of independence: if two acts agree on a sub-event, that common consequence is irrelevant to the comparison. P3 is monotonicity in outcomes (better outcomes are better in every state). P4 ensures that "which outcome wins on which event" is well-defined consistently across acts. P5–P7 are technical: non-triviality, a continuity-style condition, and the conditions needed to ensure the representation pins down a unique probability measure (rather than a finitely-additive class).

Theorem 2 (Savage representation, stated)**.** *A preference relation $\succsim$ on Savage acts satisfies P1–P7 iff there exist a unique countably additive probability measure μ on (S, Σ) and a utility $u : X \to \mathbb{R}$, unique up to positive affine transformation, such that*

$$f \succsim g \iff \int_S u(f(s))\, d\mu(s) \geq \int_S u(g(s))\, d\mu(s).$$

The proof is hard. We will not give it. The standard reference is Savage (1954); Kreps (1988) works through the structure with the level of detail this handout would aim at, and Gilboa (2009, Ch. 10) is the recent treatment.

What Savage's theorem says is a strong claim about agents: if they satisfy P1–P7, they *behave as if* they have a single probability measure on states and a utility on outcomes, regardless of whether they could write down the probability themselves. That is the formal content of the Bayesian position. The criticism, due to Ellsberg, is that the antecedent fails empirically.

3 Anscombe–Aumann acts

Savage's axiomatic lifting is heavy because Savage starts from scratch — the probability measure is to be derived, the utility on outcomes is to be derived, and both have to come out of preferences over acts that map states directly to outcomes. Anscombe and Aumann (1963) pointed out that if we are willing to assume the agent has expected-utility preferences *within each state* — that is, if we let acts map states to lotteries rather than directly to outcomes — then the axiomatic lifting becomes much lighter. The vNM machinery from the previous handout does most of the work, and the "extra" axioms beyond vNM are intuitively obvious. The framework is now the standard one for ambiguity theory, and the modern multi-prior representations in §6 and §7 are most cleanly stated in it.

Definition 3. An *Anscombe–Aumann act* (or *AA act*) is a function $f : S \to \Delta(X)$ from the state space S to the set of lotteries over an outcome space X. The agent's preferences are over AA acts.

The reading: in each state s, the agent receives the lottery $f(s) \in \Delta(X)$, which she then evaluates by expected utility under the lottery. The state determines *which* lottery she gets; the agent is uncertain over which state will obtain.

A *constant* AA act is one of the form $f(s) = p$ for every s, where $p \in \Delta(X)$ is a fixed lottery. A constant act is just a lottery; it doesn't depend on the state. The vNM theorem from the previous handout applies on constant acts directly: under the vNM axioms, the agent's preferences over constant acts admit an expected-utility representation. The AA framework's job is to extend that representation from constant acts to all acts, using a small additional set of axioms.

The axioms beyond the vNM-on-constant-acts piece:

- *Mixture independence on acts.* For any acts f, g, h and any $\alpha \in (0,1)$, $f \succsim g \iff \alpha f + (1-\alpha)h \succsim \alpha g + (1-\alpha)h$, where mixtures of acts are taken pointwise (the lottery $f(s)$ is mixed with the lottery $h(s)$, state by state).
- *Monotonicity.* If $f(s) \succsim g(s)$ as constant lotteries for every state s, then $f \succsim g$ as acts.
- *State-independence.* The agent's vNM utility on outcomes does not depend on which state is realized. (Formally: any two states are interchangeable in the sense that constant-on-each-state acts can be compared by their lottery content alone.)

The state-independence axiom is the substantive one. It rules out, e.g., a regulator whose utility for an outcome depends on which climate scenario was the source — "the same money is worth less to me in the high-warming state because I'd rather it had come from the low-warming state" is the kind of preference state-independence forbids. The axiom is plausible in many political-economy applications and implausible in others; the literature on state-dependent utility is exactly the line of work that drops it.

Theorem 4 (Anscombe–Aumann representation, stated)**.** *A preference relation $\succsim$ on AA acts satisfies the vNM axioms on constant acts, mixture independence on acts, monotonicity, and state-independence iff there exist a unique probability measure μ on (S, Σ) and a utility $u : X \to \mathbb{R}$, unique up to positive affine transformation, such that*

$$f \succsim g \iff \int_S \mathbb{E}_{f(s)}[u] \, d\mu(s) \geq \int_S \mathbb{E}_{g(s)}[u] \, d\mu(s).$$

The inner expectation $\mathbb{E}_{f(s)}[u]$ is taken under the lottery $f(s) \in \Delta(X)$ at state s, in the sense of the previous handout; the outer integral is taken under the subjective probability measure μ over states.

The structural content is the same as Savage's representation theorem: a unique probability on states, a utility on outcomes unique up to positive affine transformation, with preferences over acts determined by expected-expected-utility. The axiomatic content is much lighter, because the inner "lottery" part of an AA act inherits the vNM machinery from the previous handout for free. Kreps (1988, Ch. 13) works through the proof; Gilboa (2009, Ch. 12) gives a more recent presentation. We will use the AA framework throughout the rest of the handout because the modern multi-prior representations are stated in it.

4 The Ellsberg paradox

Ambiguity in the political-economy sense — a regulator's uncertainty over climate-sensitivity scenarios, a war planner's uncertainty over adversary type — is exactly the kind of uncertainty Savage's framework was supposed to handle by eliciting a subjective probability from preferences. The cleanest demonstration that empirical preferences resist this elicitation is due to Ellsberg (1961), in his thesis and the follow-up *Quarterly Journal of Economics* paper. The Ellsberg paradox is to choice under ambiguity what the Allais paradox is to choice under risk: a small, robust experimental finding that no single-prior expected-utility maximizer can rationalize.

The three-color version. An urn contains 90 balls. Thirty are red. The remaining sixty are some unspecified mix of black and yellow — the proportion is unknown. A ball is drawn at random. The four bets:

	Red (30)	Black (?)	Yellow (?)
Bet I	\$100	\$0	\$0
Bet II	\$0	\$100	\$0
Bet III	\$100	\$0	\$100
Bet IV	\$0	\$100	\$100

The modal experimental pattern is I $\succ$ II (the bet on the known 30/90 probability of red is preferred to the bet on the unknown probability of black) and IV $\succ$ III (the bet on the known 60/90 probability of black-or-yellow is preferred to the bet on the unknown probability of red-or-yellow). The pattern reveals that subjects prefer the known probability over the unknown one in *both* pair-comparisons, regardless of which color is the "known" one. Probability ambiguity itself is what subjects are reacting to.

Proposition 5. *No probability measure μ on $S = \{R, B, Y\}$ rationalizes both I $\succ$ II and IV $\succ$ III via expected-utility maximization.*

Proof. Suppose $u(\$100) > u(\$0)$ (which it must, for any nontrivial preference). Write μ_R, μ_B, μ_Y for the probabilities under μ. Then

- I $\succ$ II says $\mu_R \cdot u(\$100) > \mu_B \cdot u(\$100)$, hence $\mu_R > \mu_B$.
- IV $\succ$ III says $(\mu_B + \mu_Y) \cdot u(\$100) > (\mu_R + \mu_Y) \cdot u(\$100)$, hence $\mu_B + \mu_Y > \mu_R + \mu_Y$, i.e., $\mu_B > \mu_R$.

The two inequalities are contradictory. □

The Ellsberg pattern is robust. It has been replicated across populations, payoff scales, and decision-theoretic contexts, and the political-economy version — bets framed over policy outcomes, candidate-victory probabilities, or strategic environments where the relevant probability is genuinely contested — exhibits the same modal pattern.[1]

[1]The Ellsberg paradox is sometimes presented as "deeper" than the Allais paradox, and the comparison is worth taking seriously. The Allais paradox violates a separability axiom (independence) within an existing probability measure: there is a μ that gives both lotteries their probabilities, but no utility u such that expected u rationalizes

5 Knightian uncertainty

Knight (1921) drew a distinction between two epistemic situations. *Risk* is the case where the agent knows the probabilities — a coin flip with a known fair coin, a polling sample with a known sample size, a lottery ticket with a known number of winners. *Uncertainty*, in Knight's strict sense, is the case where the probabilities are unknown and possibly unknowable. Knight's examples were drawn from the entrepreneurial setting (the entrepreneur's profit comes from bearing genuine uncertainty about the market, not from bearing measurable risk that could be insured), but the distinction is general.

The Bayesian programme, taking shape over the twenty years following Knight's book, argued that the distinction was illusory. Probabilities, on the Bayesian view, are subjective and can always be elicited from preferences. The agent who claims not to have probabilities is, the Bayesians argued, simply not introspecting carefully enough. Savage (1954)'s axiomatic vindication of this view is the pinnacle of the programme. Risk and uncertainty are the same kind of thing; the agent's task is to articulate the probability, and her preferences are the diagnostic tool that lets the analyst extract it.

Ellsberg's experimental finding cuts against this. The Knightian view that the modern literature has converged on is that the agent's epistemic state, in many real situations, is a *set* of probability measures rather than a single one. The set may be wide (no information about the distribution of unknown balls) or narrow (some information, but not enough to pin down a unique measure). The agent's choice rule has to take the set into account. The most influential proposal is the maxmin criterion of the next section.

The political-economy applications where the Knightian view has the most bite are the ones where model uncertainty is the live concern: a regulator who has several scientific models of climate sensitivity but no agreed-upon prior over them; a central banker with several macroeconomic models of how a policy will transmit; a war planner with several models of an adversary's type. In each case, asking the decision-maker to "elicit a subjective probability" over the models is asking her to pretend confidence she does not have, and the maxmin criterion is one principled way of taking the lack of confidence seriously.

6 Gilboa–Schmeidler maxmin expected utility

The centerpiece of the modern ambiguity literature is the maxmin expected utility model of Gilboa and Schmeidler (1989). The setting is the Anscombe–Aumann framework. The agent has a closed convex set of priors $\mathcal{P} \subseteq \Delta(S)$ rather than a single μ, and she evaluates an act by the worst-case expected utility over the set:

$$V(f) = \min_{\mu \in \mathcal{P}} \int_S \mathbb{E}_{f(s)}[u]\, d\mu(s).$$

the modal preference pattern. The Ellsberg paradox violates the existence-of-a-single-probability-measure axiom: not only is there no u that rationalizes the pattern, there is no μ either. Allais says expected-utility maximization fails; Ellsberg says *subjective expected utility* (the joint thesis that one has a probability and a utility) fails. The Knightian distinction — between risk (known probabilities) and uncertainty (unknown probabilities) — was sidelined for half a century by the Bayesian programme, on the basis that any honest agent's uncertainty admits a subjective probability via Ramsey–de-Finetti–Savage elicitation. Ellsberg (1961) is what reopened the question. Knight (1921) and Keynes (1921) both made the case for a non-probabilistic conception of uncertainty in the 1920s, in close proximity to each other; both books are worth reading on the conceptual question of what *ambiguity* actually is, and they remain the philosophical reference for the modern formal literature on multi-prior decision rules.

Her preference is then $f \succsim g \iff V(f) \geq V(g)$. The minimum is well-defined because $\mathcal{P}$ is closed and convex and the inner expression is continuous in μ.

The substantive content of the model is that the agent hedges against the worst plausible prior in $\mathcal{P}$. "Plausible" is encoded entirely by what is in $\mathcal{P}$: priors outside it are excluded as implausible (by the agent's evidence or by the analyst's modeling assumption), and priors inside it are all considered live. The agent does not aggregate them by an outer probability or a smooth function; she takes the worst case.

The axiomatic content is the modification of the AA axioms that produces this representation. Gilboa and Schmeidler (1989) weaken mixture independence to *certainty-independence* (independence holds when the third act is constant) and add an *uncertainty-aversion* axiom: a convex combination of two indifferent acts is weakly preferred to either.

Theorem 6 (Gilboa–Schmeidler representation, stated)**.** *A preference relation $\succsim$ on AA acts satisfies the vNM axioms on constant acts, monotonicity, state-independence, certainty-independence, and uncertainty aversion iff there exist a unique non-empty closed convex set $\mathcal{P} \subseteq \Delta(S)$ and a utility $u : X \to \mathbb{R}$, unique up to positive affine transformation, such that*

$$f \succsim g \iff \min_{\mu \in \mathcal{P}} \int_S \mathbb{E}_{f(s)}[u]\, d\mu(s) \geq \min_{\mu \in \mathcal{P}} \int_S \mathbb{E}_{g(s)}[u]\, d\mu(s).$$

The proof is in Gilboa and Schmeidler (1989); Gilboa (2009, Ch. 13) works through the structure. We illustrate by computing the Ellsberg pattern under maxmin EU.

Example 7 (Ellsberg under maxmin EU)**.** Take $S = \{R, B, Y\}$, the three-color Ellsberg states. The agent knows $\mu(R) = 30/90 = 1/3$ but is uncertain about $(\mu(B), \mu(Y))$ subject to $\mu(B) + \mu(Y) = 2/3$. Her prior set is

$$\mathcal{P} = \big\{(1/3,\ p,\ 2/3 - p) : p \in [0, 2/3]\big\}.$$

Take $u(\$100) = 1$ and $u(\$0) = 0$. The maxmin values:

$$\begin{aligned}
V(\mathrm{I}) &= \min_{p \in [0,2/3]} \big[1/3 \cdot 1 + p \cdot 0 + (2/3 - p) \cdot 0\big] = 1/3, \\
V(\mathrm{II}) &= \min_{p \in [0,2/3]} \big[1/3 \cdot 0 + p \cdot 1 + (2/3 - p) \cdot 0\big] = 0, \\
V(\mathrm{III}) &= \min_{p \in [0,2/3]} \big[1/3 \cdot 1 + p \cdot 0 + (2/3 - p) \cdot 1\big] = 1/3, \\
V(\mathrm{IV}) &= \min_{p \in [0,2/3]} \big[1/3 \cdot 0 + p \cdot 1 + (2/3 - p) \cdot 1\big] = 2/3.
\end{aligned}$$

The minimum in (II) is at $p = 0$ (the worst case for a bet on black: there are no black balls). The minimum in (III) is at $p = 2/3$ (the worst case for a bet on red-or-yellow: all the unknown balls are black, so yellow contributes nothing). The bets in (I) and (IV) are bets on *unions of events with known total probability*, so the inner expression is constant in p and the minimum is just the constant value.

So under maxmin EU: $V(\mathrm{I}) = 1/3 > 0 = V(\mathrm{II})$ and $V(\mathrm{IV}) = 2/3 > 1/3 = V(\mathrm{III})$. Both modal preferences are rationalized.

The example displays the structural content of the maxmin model. The bets on known events — "red," which has known probability 1/3, and "black-or-yellow," which has known probability 2/3 — have constant value across the prior set, because the unknown p doesn't affect them. The bets on unknown events — "black" alone, or "red-or-yellow" — are evaluated at the worst case of p, which gives them lower value than a single-prior calculation at $p = 1/3$ would. Ambiguity aversion shows up in the fact that bets on unknown events get penalized by the worst-case minimum.[2]

Example 8 (Climate-policy regulator)**.** A regulator must choose between two policies. The states are $S = \{H, L\}$ (high climate sensitivity, low climate sensitivity), with the regulator unsure of the prior over them: $\mathcal{P} = \{(\mu_H, 1 - \mu_H) : \mu_H \in [0.3, 0.7]\}$, capturing scientific disagreement that doesn't pin down a unique prior. Policy A is a strong mitigation rule: it costs $-c$ in either state but reduces damages, giving net utility $u_H = 5$ in H and $u_L = 1$ in L (the cost is justified in H, less so in L). Policy B is a weak mitigation rule: net utility $u_H = 0$ in H (under-mitigation, large damages) and $u_L = 4$ in L (cost-effective). The maxmin values: $V(A) = \min_{\mu_H \in [0.3,0.7]}(5\mu_H + 1(1 - \mu_H)) = \min(4\mu_H + 1) = 0.3 \cdot 4 + 1 = 2.2$, attained at $\mu_H = 0.3$ (the worst-case for a strong-mitigation policy is low sensitivity). $V(B) = \min_{\mu_H \in [0.3,0.7]}(0\mu_H + 4(1 - \mu_H)) = \min(-4\mu_H + 4) = 4 - 4 \cdot 0.7 = 1.2$, attained at $\mu_H = 0.7$ (the worst-case for a weak-mitigation policy is high sensitivity). So $V(A) > V(B)$: the regulator chooses strong mitigation. Under SEU with any single prior in $\mathcal{P}$, the choice would depend on μ_H (specifically, $A \succ B$ iff $5\mu_H + 1(1 - \mu_H) > 4(1 - \mu_H)$, i.e., $\mu_H > 3/8 = 0.375$). The maxmin criterion picks the strong mitigation policy because, at *every* μ_H in the prior set, strong mitigation does at least adequately, while weak mitigation does poorly under high sensitivity.

7 Choquet expected utility

The Gilboa–Schmeidler model is one way to extend the AA representation to ambiguity. Schmeidler (1989) gave a closely related extension, also stated in the AA framework, that uses a different mathematical object on the state space. Where Gilboa–Schmeidler replaces the single probability measure with a closed convex set of measures, Schmeidler replaces it with a single *capacity* — a non-additive set function on the state space. The two extensions are not duplicates: every closed convex set of priors gives rise to a capacity (its lower envelope), and every *convex* capacity gives rise to a set of priors (its core), and on convex capacities the two models give the same evaluation of acts. Outside the convex case the two diverge.

Definition 9. A *capacity* on (S, Σ) is a function $\nu : \Sigma \to [0, 1]$ with $\nu(\emptyset) = 0$, $\nu(S) = 1$, and *monotonicity*: $A \subseteq B \Rightarrow \nu(A) \leq \nu(B)$. The capacity is *convex* (or *supermodular*) if $\nu(A \cup B) + \nu(A \cap B) \geq \nu(A) + \nu(B)$ for all $A, B \in \Sigma$.

A capacity is a probability measure with the additivity axiom dropped. The probability measure is the special case where $\nu(A \cup B) = \nu(A) + \nu(B)$ for disjoint A, B. Convex capacities are

[2] The substantive new axiom in Gilboa–Schmeidler beyond the AA framework is *uncertainty aversion*: a convex combination of two indifferent acts is weakly preferred to either. The reading: hedging is good. Two indifferent acts may have non-overlapping "best states" (states where each one delivers high outcomes), and a convex mix smooths across both, reducing exposure to the worst case under any single prior. The set-of-priors representation is the formal content of that hedging instinct: the worst-case prior is what the agent hedges against, and a mix of two acts performs at-least-as-well under *every* prior in $\mathcal{P}$ as the worse of the two on its own. The connection to robust control and robust decision-making is direct. Hansen and Sargent (2008) develop an entire macroeconomic-policy framework around the maxmin criterion under the heading *robustness*: the policymaker designs a policy that performs adequately under the worst plausible model in a class of models, rather than optimally under a single model. The political-economy applications where this framework has the most traction are exactly the ones where model-class uncertainty is the live concern: regulatory policy under contested science, monetary policy under structural uncertainty, deterrence under uncertainty about adversary type. In each, the maxmin criterion is a principled formalization of "make the policy robust to which model turns out to be right," and Gilboa–Schmeidler is the decision-theoretic foundation that licenses it.

between probability measures and arbitrary capacities; the convexity inequality is what produces the connection to the maxmin model.

Definition 10. The *Choquet integral* of a function $\varphi : S \to \mathbb{R}$ with respect to a capacity ν is

$$\int_S \varphi \, d\nu = \int_0^\infty \nu(\{s : \varphi(s) \geq t\}) \, dt + \int_{-\infty}^0 \big[\nu(\{s : \varphi(s) \geq t\}) - 1\big] \, dt,$$

defined whenever both Riemann integrals exist (and reducing to the second integral being zero when $\varphi \geq 0$).

The Choquet integral is the cumulative-distribution-style integral that makes sense without additivity. For an additive capacity (i.e., a probability measure), the Choquet integral coincides with the standard Lebesgue expectation, so the Choquet integral generalizes expectation in the way capacities generalize probability measures.

Theorem 11 (Schmeidler representation, stated)**.** *A preference relation $\succsim$ on AA acts satisfies the vNM axioms on constant acts, monotonicity, state-independence, and* comonotonic-independence *(mixture independence holds when both pairs of acts are comonotonic — ranked the same way across states) iff there exist a unique capacity ν on (S, Σ) and a utility $u : X \to \mathbb{R}$, unique up to positive affine transformation, such that*

$$f \succsim g \iff \int_S \mathbb{E}_{f(s)}[u] \, d\nu \geq \int_S \mathbb{E}_{g(s)}[u] \, d\nu,$$

where the integrals are Choquet integrals.

The connection to Gilboa–Schmeidler. For a convex capacity ν, define the *core* of ν as

$$\text{core}(\nu) = \{\mu \in \Delta(S) : \mu(A) \geq \nu(A) \text{ for every } A \in \Sigma\}.$$

Theorem 12 (Choquet–maxmin equivalence on convex capacities)**.** *If ν is a convex capacity, then* core(ν) *is non-empty, closed, convex, and*

$$\int_S \varphi \, d\nu = \min_{\mu \in \text{core}(\nu)} \int_S \varphi \, d\mu$$

for every bounded measurable φ. Consequently, Choquet EU under a convex capacity ν coincides with maxmin EU under prior set $\mathcal{P} = \text{core}(\nu)$.

The result ties the two models together exactly. A convex capacity is the same data as its core. Choquet EU under the capacity is maxmin EU over the core. Outside the convex case, Choquet EU is more general (the capacity may have empty core, and the Choquet representation still works), but at the cost of giving up the maxmin reading. Gilboa (2009, Ch. 14) works through the equivalence and the divergence in the non-convex case.

8 What's next

The decision-theory cluster closes here. Three strands extend the material:

- *Game theory under uncertainty.* The state space in this handout was the source of the agent's uncertainty; in a game-theoretic setting the relevant uncertainty includes other players' strategies, and the state space is partly endogenous. The Bayesian and ambiguity-theoretic apparatus extends, with care, to game-theoretic settings; the literature on Bayesian games (Harsanyi 1967–1968) and on games under ambiguity is the natural successor reading.

- *Dynamic consistency under ambiguity.* The maxmin criterion over a fixed prior set has been the static story. Once the agent updates her prior set on the basis of new information, dynamic-consistency questions arise: is the updated $\mathcal{P}$ what the agent would have committed to at the outset? The answer is sometimes no, and the formal literature on dynamic ambiguity (Epstein–Schneider, among others) takes this seriously.

- *Other multi-prior models.* This handout featured Gilboa–Schmeidler maxmin EU and Schmeidler's Choquet EU. The literature has additional families: *variational preferences* (Maccheroni, Marinacci, and Rustichini, 2006, replacing the maxmin criterion with $\min_\mu \left[\int u\, d\mu + c(\mu)\right]$, where c is a penalty function on priors), and *smooth ambiguity* (Klibanoff, Marinacci, and Mukerji, 2005, replacing the maxmin with a second-order expectation that smooths across priors). Both extend the framework but at the cost of additional axiomatic and parametric structure that this handout does not pursue.

For the standard graduate references, Savage (1954) and Kreps (1988) for the Savage framework, Anscombe and Aumann (1963) for the AA framework (the original five-page paper is itself worth reading), Gilboa and Schmeidler (1989) and Schmeidler (1989) for the source theorems of §6 and §7, Gilboa (2009) for the unified treatment, and Hansen and Sargent (2008) for the macroeconomic-policy applications under the robustness heading. For political-economy applications specifically, Austen-Smith and Banks (1999) and Austen-Smith and Banks (2005) address Bayesian decision theory throughout; the modern political-economy literature on ambiguity is more scattered, with applications in regulatory policy, deterrence theory, and constitutional design appearing under various headings.

9 Exercises

Exercise 13. Verify the Anscombe–Aumann representation theorem on a small example. Let $S = \{s_1, s_2\}$ and $X = \{x_1, x_2\}$ with $u(x_1) = 0, u(x_2) = 1$, and $\mu(s_1) = 0.6, \mu(s_2) = 0.4$. Consider the AA acts f (delivering lottery $0.5x_1 + 0.5x_2$ in both states) and g (delivering x_2 in s_1 and x_1 in s_2). Compute $V(f), V(g)$ and identify which act the agent prefers.

Exercise 14. Show that Savage's representation, restricted to constant acts, is the same as the vNM theorem from the previous handout, with the lottery $f(s) = p$ identified with $p \in \Delta(X)$. (Hint: a constant Savage act is just an outcome; a constant AA act is a lottery; the AA representation on constant acts is the vNM representation.)

Exercise 15. Verify Proposition 5 (Ellsberg paradox) directly: assume an SEU representation exists, suppose both I $\succ$ II and IV $\succ$ III hold, and derive the contradiction $\mu(R) > \mu(B) > \mu(R)$.

Exercise 16. *Climate-policy regulator under model uncertainty.* Continuing Example 8: suppose a third policy C is added with utility $u_H = 3, u_L = 3$ (a cost-balanced policy that performs identically across states). Compute $V(A), V(B), V(C)$ under the prior set $\mathcal{P} = \{(\mu_H, 1 - \mu_H) : \mu_H \in [0.3, 0.7]\}$. Which policy does the maxmin criterion select? Discuss in one or two sentences whether the constant-utility policy C is what the regulator should always choose under maxmin reasoning.

Exercise 17. *Deterrence under Knightian uncertainty.* A war planner contemplates an action that produces outcome x_H if the adversary's resolve type is high (he stands firm) and outcome x_L if it is low (he backs down). The planner has no frequency-based prior over resolve type and uses the maxmin criterion with $\mathcal{P} = \{(\mu_H, 1 - \mu_H) : \mu_H \in [0.2, 0.8]\}$. Suppose $u(x_L) = 1$ (the planner gets her way) and $u(x_H) = -3$ (an unwinnable confrontation). Compute the maxmin value of taking the action versus not taking it (which gives certain $u = 0$). For what range of $u(x_L)$ values is the action preferred? Discuss in one sentence the "ambiguity premium" on resolve uncertainty: the additional caution induced by the prior set being wider than a single μ_H.

Exercise 18. *Constitutional design as choice under structural uncertainty.* A constitutional designer chooses between two voting rules: R (simple majority) and S (supermajority). The relevant states are {*cooperative coalitions form*, *polarized coalitions form*} with prior set $\mathcal{P}$ allowing wide variation. Under cooperative coalitions, R delivers $u = 4$ and S delivers $u = 3$ (the supermajority is unnecessarily restrictive). Under polarized coalitions, R delivers $u = 0$ (majority tyranny) and S delivers $u = 2$. (a) For what prior set $\mathcal{P}$ does maxmin EU select S? (b) Discuss in two or three sentences how the maxmin reading of this scenario lines up with Rawls's veil-of-ignorance argument for institutional choice, and how it differs from the Harsanyi-EU reading of the same scenario.

Exercise 19. Let $\mathcal{P}$ be the singleton $\{\mu_0\}$. Show that maxmin EU under $\mathcal{P}$ coincides with SEU under μ_0. Conclude that SEU is the special case of Gilboa–Schmeidler where the prior set has collapsed to a single measure — ambiguity disappears when there is no genuine multiplicity of priors.

Exercise 20. Verify the Choquet–maxmin equivalence on a small example. Let $S = \{s_1, s_2, s_3\}$ and define a capacity ν by $\nu(\{s_i\}) = 1/4$ for each i, $\nu(\{s_i, s_j\}) = 1/2$ for each pair, $\nu(S) = 1$, $\nu(\emptyset) = 0$. (a) Verify that ν is convex. (b) Compute the core. (c) Take a function $\varphi(s_1) = 1, \varphi(s_2) = 2, \varphi(s_3) = 3$ and verify that $\int \varphi \, d\nu = \min_{\mu \in \mathrm{core}(\nu)} \int \varphi \, d\mu$.

Exercise 21. Show that comonotonic-independence (the Schmeidler axiom) is implied by full mixture independence (the AA axiom) but is strictly weaker. Give an example of two acts that are not comonotonic, and explain in one sentence why mixing them with a third act might fail mixture independence but trivially satisfy comonotonic-independence.

Exercise 22. Reflect on the cluster as a whole. The three handouts have moved from known alternatives (#18) through known probabilities (#19) to unknown probabilities (#20). Identify, in two or three sentences, where in this progression the political-economy applications you encounter most often actually live — and which formal framework (rationalizability, expected utility, maxmin EU) is the one you find yourself reaching for in practice.

References

Anscombe, F. J. and R. J. Aumann (1963). "A Definition of Subjective Probability". In: *Annals of Mathematical Statistics* 34.1, pp. 199–205.

Austen-Smith, David and Jeffrey S. Banks (1999). *Positive Political Theory I: Collective Preference*. Ann Arbor: University of Michigan Press.

– (2005). *Positive Political Theory II: Strategy and Structure*. Ann Arbor: University of Michigan Press.

Ellsberg, Daniel (1961). "Risk, Ambiguity, and the Savage Axioms". In: *Quarterly Journal of Economics* 75.4, pp. 643–669.

Gilboa, Itzhak (2009). *Theory of Decision under Uncertainty*. Cambridge: Cambridge University Press.

Gilboa, Itzhak and David Schmeidler (1989). “Maxmin Expected Utility with a Non-Unique Prior”. In: *Journal of Mathematical Economics* 18.2, pp. 141–153.

Hansen, Lars Peter and Thomas J. Sargent (2008). *Robustness*. Princeton: Princeton University Press.

Keynes, John Maynard (1921). *A Treatise on Probability*. London: Macmillan.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005). “A Smooth Model of Decision Making under Ambiguity”. In: *Econometrica* 73.6, pp. 1849–1892.

Knight, Frank H. (1921). *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin.

Kreps, David M. (1988). *Notes on the Theory of Choice*. Boulder, CO: Westview Press.

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini (2006). “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences”. In: *Econometrica* 74.6, pp. 1447–1498.

Savage, Leonard J. (1954). *The Foundations of Statistics*. New York: Wiley.

Schmeidler, David (1989). “Subjective Probability and Expected Utility without Additivity”. In: *Econometrica* 57.3, pp. 571–587.