Game Theory

Decision–Theoretic Foundations

Lasha Chochua

2026

What Is Game Theory?

Game Theory – Definition

Definition: Game Theory

The study of mathematical models of **conflict and cooperation** between **intelligent rational** decision-makers.

- Analyzes situations where two or more individuals make decisions that **influence one another's welfare**
- Provides general mathematical techniques for such interdependent situations
- Alternative names: "conflict analysis," "interactive decision theory"
- Offers insights of fundamental importance across **all branches** of the social sciences and for practical decision-makers

Game Theory – Historical Roots

- **Zermelo (1913):** Chess has a determined outcome under perfect play
- **Borel (1921):** First formal treatment of mixed strategies
- **von Neumann (1928):** Minimax theorem for zero-sum games
- **von Neumann & Morgenstern (1944):** *Theory of Games and Economic Behavior* – the foundational text

- Much early work at Princeton during WWII, alongside leaders of theoretical physics
- Motivation: physical sciences created a nuclear dilemma – game theory offers tools to understand **social systems** for moderating human conflict in response

Rationality – Definition

Definition: Rationality

A decision-maker is **rational** if he makes decisions **consistently in pursuit of his own objectives** – formally, he maximizes the expected value of his payoff measured in some utility scale.

- This is a **technical** definition, not colloquial usage
- Justified by the von Neumann–Morgenstern (1947) expected-utility theorem
- Based on remarkably **weak consistency assumptions** – the axioms we will study
- Rationality alone does not require understanding the full strategic structure of the situation

Intelligence – Definition

Definition: Intelligence

A player is **intelligent** if he knows **everything that we (the analysts) know** about the game and can make any inferences we can make.

- If our theory correctly predicts behavior, **each player also understands that prediction**
- A good theory must survive being publicly known
- Contrast with **price theory**: agents are rational but only perceive price signals – they do not model the whole economic structure and are not assumed to be intelligent in this sense

Rationality + Intelligence – The Strategic Circularity

- Player 1 tries to predict Player 2's choice
- To do so, Player 1 imagines herself in Player 2's position
- Player 2 is doing exactly the same – anticipating Player 1's reasoning

- Each player's optimal decision **depends on solving the other player's problem**
- Neither problem can be solved without understanding the solution to the other
- The two decision problems must be **analyzed together, like a system of equations**

> **Note**
>
> This is precisely why single-agent decision theory is insufficient for strategic settings – and why game theory is needed as an independent discipline.

Why Self-Consistency Matters

- If a theory predicts that some individuals will be **systematically fooled** or led into making costly mistakes, the theory will tend to **lose its validity**
- Learning can come from experience or from reading the published theory itself

- Example: if a theory predicts Player 2 always plays Left, then Player 1 will exploit this
- But then Player 2, knowing the theory, will not play Left – the theory contradicts itself

Note

The importance of game theory in the social sciences is largely derived from this self-consistency requirement: **a good theory must survive being known by the players it describes.**

Game Theory in Economics

From Niche to Mainstream

- In the **1960s–70s**, game theory occupied an isolated niche – pursued by specialists, unknown to most economists
- By the **1980s–90s**, a burst of foundational work (Rubinstein 1982; Kreps–Wilson 1982; Milgrom–Roberts 1982) triggered explosive adoption
- Today: game theory is a **standard tool** across all fields of economics

- The transition is visible in the first-year graduate theory sequence:
  - Old structure: optimization → competitive equilibrium
  - New structure: optimization → **game theory**

Note

Game theory did not merely occupy a bridgehead – it extended its reach throughout economics until familiarity replaced excitement (Samuelson 2016).

Game Theory and Methodological Individualism

- Economics is built on **methodological individualism**: explain social phenomena from individual behavior. The two pillars of this program:

First Pillar	Second Pillar (original)
Individual optimization: utility and profit maximization	Competitive markets: price-taking equilibrium
	Second Pillar (modern)
	Game theory: strategic equilibrium

- Competitive markets = limiting case of Cournot as the number of firms $\to \infty$ – one way to nest the two frameworks
- This integration is not universal: the **Chicago price theory** tradition retained competitive markets as the primary tool and was skeptical of game-theoretic modeling as the general case

Classical vs. Instrumental View

Classical View	Instrumental View
The game is a **literal, complete description** of the situation	The game is a **deliberate approximation** designed to capture what matters
Equilibrium is deducible from common knowledge of rationality alone	Equilibrium concept is part of model construction, informed by context
Cournot vs. Bertrand: check what firms *actually* do	Cournot vs. Bertrand: use whichever gives the most useful insight
More realism $\Rightarrow$ better model	More realism $\not\Rightarrow$ better model

- The failure of the **equilibrium refinements program** – which sought a single canonical solution concept – pushed the field firmly toward the instrumental view

The Multiple Equilibria Problem

- Many games have **multiple Nash equilibria** – perhaps the central challenge
- Repeated games are especially rich: if players are sufficiently patient, the **folk theorem** (Fudenberg–Maskin 1986) says virtually any outcome is an equilibrium
- Three responses:
  - Focus on results that hold for **all** equilibria (like welfare theorems in competitive theory)
  - Let **empirical methods** select the equilibrium – assume data reflect some consistent equilibrium
  - **Embrace multiplicity** when the application calls for it

Note

Schelling (1960): **focal points** – context, history, and cultural salience guide players toward one equilibrium without being in the formal model. This is instrumental game theory at work.

Where Game Theory Succeeded – and Where It Struggled

Application	Success?	Why?
Auctions	**Yes** – transformed resource allocation worldwide	Theory isolated the first-order tradeoff (bid shading vs. winning probability); intuition is reliable
Matching	**Yes** – reshaped labor markets, kidney exchange, school assignment	Cooperative + noncooperative foundations both available
Bargaining	**Partial** – rich theory but limited practical impact	Outcomes too sensitive to fine modeling details (timing, horizon, who moves first)

- Progress in game theory comes less from the **science** of applying existing models and more from the **art** of formulating the right ones

Cooperative vs. Noncooperative Game Theory

- Early game theory (von Neumann–Morgenstern 1944) was dominated by **cooperative** approaches
- The field later swung hard toward **noncooperative** theory – today's first-year courses may not mention the core at all

- Key cooperative solution concepts:
  - **Core:** no coalition can break away and do better for all its members
  - **Shapley value:** allocates payoffs proportional to marginal contributions to coalitions
  - **Nucleolus:** minimizes the maximum dissatisfaction of any coalition

- The **Nash program** (Nash 1953): find noncooperative foundations for cooperative solutions – and cooperative characterizations of noncooperative outcomes
- **Matching theory** is the leading modern synthesis: stable matchings (cooperative) implemented by the Gale–Shapley deferred acceptance algorithm (noncooperative)

Game Theory Beyond Economics (I)

- Game theory is now the **language of the social sciences** – and increasingly beyond
- Applications span:
  - **Political science:** legislative behavior, lobbying, electoral competition, international relations (Ostrom 1990 on commons governance)
  - **Biology:** evolutionary game theory – perhaps the greatest empirical success (Maynard Smith–Price 1973)
  - **Law, philosophy, neuroscience, computer science, engineering**

Game Theory Beyond Economics (II)

- The breadth of "players" is striking: people, firms, countries, political parties, cells, neurons, animals, routers
- Many of these interpretations are incompatible with a **classical rational actor** – which is precisely why the instrumental view dominates as game theory spreads

Note

The success of evolutionary game theory – selecting Nash equilibria as stable rest points of adaptive dynamics – provides the **dynamic foundation** that classical game theory lacked.

What This Means for This Course (I)

- We follow the **instrumental approach**: game-theoretic models are tools, judged by the quality of their insights
- Our foundation is **decision theory**: rationality, utility, beliefs
  - Before we can analyze strategic interaction, we need a precise model of individual choice
  - Choice axioms pin down exactly what "rational" means – and what it does not
- Then we build up:
  - Games in normal and extensive form $\rightarrow$ Nash equilibrium $\rightarrow$ refinements $\rightarrow$ repeated games (if time allows)

What This Means for This Course (II)

Note

The self-consistency requirement – a good theory must survive being known by the players – is why game theory needs its own foundations, distinct from single-agent decision theory. That is where we begin.

Decision Theory Foundations

What Are We Building? – The Big Picture (I)

Goal

Construct a model in which the behavior of any rational, intelligent agent can be described by **expected utility maximization** with **subjective beliefs** – and show this follows from remarkably weak consistency requirements.

- The central object of choice is a **lottery**: a function mapping states of the world to probability distributions over prizes
- We impose consistency axioms on preferences over lotteries
- We then prove that: preferences satisfying the axioms can be represented as maximization of $E_p[u(f) \mid S]$
- The **two key ingredients** elicited from preferences: utility $u(x, t)$ and beliefs $p(t \mid S)$

What Are We Building? – The Big Picture (II)

Road map

Formal setting → 8 axioms → Theorem 1 (existence) → Theorems 2–3 (uniqueness) → behavioral paradoxes → domination

The Object of Choice – Intuition First

- A **prize** $x \in X$ is everything the agent cares about: money, consumption, outcomes
- A **state** $t \in \Omega$ is a subjective unknown – something nature picks that the agent does not control

- **Example.** A firm decides whether to enter a new market.
  - Prizes: profit outcomes ($\$0$, $\$1M$, $\$5M$)
  - States: competitor enters or does not; demand is high or low
  - The firm does not know the state when it decides
  - But even after the state is realized, there may be residual **objective** randomness (e.g., a government contract awarded by lottery)

- The agent's choice is a **lottery** f: for each state t, a probability distribution $f(\cdot \mid t)$ over prizes
- This single object captures **both** objective risk (known probabilities) **and** subjective uncertainty (unknown states)

The Formal Setting – Prizes, States, Distributions

- X = finite set of possible **prizes**
  - A prize is a **complete specification** of all aspects the decision-maker cares about
  - Prizes are mutually exclusive and exhaustive
- Ω = finite set of possible **states of the world**
  - One state will be the true state
  - A state summarizes all **subjective unknowns** that might affect the prize received
- For any finite set Z, define the simplex of probability distributions:

$$\Delta(Z) = \left\{ q : Z \to \mathbb{R} \,\middle|\, \sum_{y \in Z} q(y) = 1 \text{ and } q(z) \geq 0,\ \forall z \in Z \right\} \tag{1}$$

Two Models of Uncertainty

Risk (Objective)	Uncertainty (Subjective)
Probabilities are **known** and agreed upon	Probabilities are **unknown** – no objective basis
"Risk" (Knight 1921)	"Uncertainty" (Knight 1921)
"Roulette lotteries" (Anscombe–Aumann): spin of a wheel, coin toss	"Horse lotteries" (Anscombe–Aumann): bet on which horse wins
Coin tosses, roulette wheels	Sports results, stock market, competitor's strategy

- The **Anscombe–Aumann** framework nests both: an act f maps states (subjective uncertainty) to roulette lotteries (objective risk)
- So $f : \Omega \to \Delta(X)$ – the horse lottery's outcome is itself a roulette lottery
- This compound structure lets the theorem separately identify beliefs p over Ω and utility u over X

Lotteries – Definition (I)

Definition: Lottery

A **lottery** is any function $f : \Omega \to \Delta(X)$, i.e. $f \in L$ where:

$$L = \{f : \Omega \to \Delta(X)\}$$

The number $f(x \mid t)$ is the **objective conditional probability** of prize x given that state t is the true state.

- $f(\cdot \mid t) \in \Delta(X)$: the prize distribution in state t – pure objective risk once t is fixed
- L is our universal object of choice – the agent picks a function, not just a number

Lotteries – Definition (II)

Interpretation

Nature first picks state t (subjective uncertainty); then the lottery $f(\cdot \mid t)$ determines the prize by objective randomization (risk). The agent chooses f before either is resolved.

Lotteries – Definition (III)

Example. $\Omega = \{t_H, t_L\}$, $X = \{\$0, \$1M, \$5M\}$

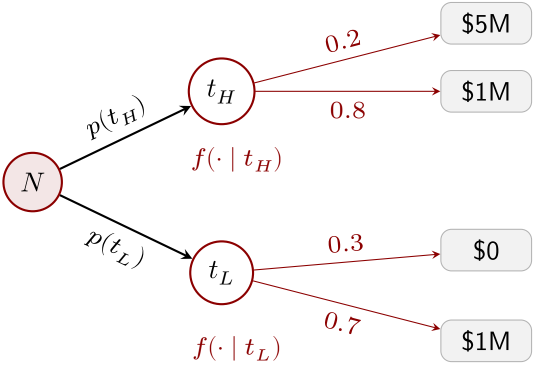

- **Left arrows** $(N \to t)$: subjective – agent does not know $p(t_H)$
- **Right arrows** $(t \to x)$: objective – probabilities $f(x \mid t)$ are known

Lotteries – Definition (IV): Matrix Representation

- An act $f \in L$ with $\Omega = \{t_1, t_2, t_3\}$ and $X = \{x_1, x_2, x_3\}$ can be written as a matrix:

$$f = \begin{bmatrix} & x_1 & x_2 & x_3 \\ t_1 & f(x_1 \mid t_1) & f(x_2 \mid t_1) & f(x_3 \mid t_1) \\ t_2 & f(x_1 \mid t_2) & f(x_2 \mid t_2) & f(x_3 \mid t_2) \\ t_3 & f(x_1 \mid t_3) & f(x_2 \mid t_3) & f(x_3 \mid t_3) \end{bmatrix}$$

- Each **row** sums to 1: it is the prize lottery conditional on that state
- Each **column** is indexed by a prize; the entry is its objective probability given the state

Lotteries – Definition (V)

- **Example** (same lottery as the tree diagram):

$$f = \begin{bmatrix} & \$0 & \$1M & \$5M \\ t_H & 0 & 0.8 & 0.2 \\ t_L & 0.3 & 0.7 & 0 \end{bmatrix}$$

Note

The matrix is fully equivalent to the tree: it just suppresses the unknown left-branch probabilities $p(t_H)$, $p(t_L)$ – those are what the theory recovers from preferences.

Compound and Degenerate Lotteries

- For $\alpha \in [0,1]$ and $f, g \in L$, the **compound lottery** $\alpha f + (1-\alpha)g$ satisfies:

$$(\alpha f + (1-\alpha)g)(x \mid t) = \alpha f(x \mid t) + (1-\alpha)g(x \mid t), \quad \forall x \in X,\ \forall t \in \Omega$$

- Interpretation: draw black ball (prob α) $\Rightarrow$ play f; draw white ball (prob $1-\alpha$) $\Rightarrow$ play g

- For any prize $x \in X$, the **degenerate lottery** $[x]$ gives prize x for sure:

$$[x](y \mid t) = \begin{cases} 1 & \text{if } y = x \\ 0 & \text{if } y \neq x \end{cases}$$

- So $\alpha[x] + (1-\alpha)[y]$ gives prize x with prob α and prize y with prob $1-\alpha$, regardless of state

Events and Conditional Preferences

- An **event** S is any nonempty subset of Ω
- $\Xi = \{S \mid S \subseteq \Omega,\ S \neq \emptyset\}$ = collection of all events

- $f \succsim_S g$: lottery f is at least as desirable as g, given the true state is in S
- Derived relations:
  - $f \sim_S g$ iff $f \succsim_S g$ and $g \succsim_S f$ (indifference given S)
  - $f \succ_S g$ iff $f \succsim_S g$ and $g \not\succsim_S f$ (strict preference given S)

- Write $\succsim, \succ, \sim$ for $\succsim_\Omega, \succ_\Omega, \sim_\Omega$ – prior preferences before any observation

- We require well-defined preferences conditional on **any** event, including zero-probability events

The Axioms

Scope of the Axioms

- All axioms hold for all $e, f, g, h \in L$, events $S, T \in \Xi$, and $\alpha, \beta \in [0, 1]$

- The axioms characterize **rational preferences** over lotteries
- They are **weak consistency assumptions** – not strong behavioral restrictions
- Together with Theorem 1, they imply expected-utility maximization

- We present 8 axioms:
  - **1A/B:** Completeness and transitivity
  - **2:** Relevance
  - **3:** Monotonicity
  - **4:** Continuity
  - **5A/B:** Objective substitution (independence)
  - **6A/B:** Subjective substitution (sure-thing principle)
  - **7:** Interest
  - **8:** State neutrality (optional)

Axiom 1A – Completeness

Axiom 1A (Completeness)

For any event S and any lotteries $f, g \in L$:

$$f \succsim_S g \quad \text{or} \quad g \succsim_S f$$

- The decision-maker can **always compare** any two lotteries
- He is never "unable to decide" in a way that leaves the relation undefined
- Note: both can hold simultaneously – that is indifference $f \sim_S g$

Axiom 1B – Transitivity

Axiom 1B (Transitivity)

For any event S and any lotteries $f, g, h \in L$:

$$\text{If } f \succsim_S g \text{ and } g \succsim_S h, \text{ then } f \succsim_S h.$$

- Preferences are **consistent** across comparisons
- Together with 1A: preferences form a **complete transitive order** over L

- Axiom 1B implies further transitivity results:
  - If $f \sim_S g$ and $g \sim_S h$ then $f \sim_S h$
  - If $f \succ_S g$ and $g \succsim_S h$ then $f \succ_S h$

Axiom 2 – Relevance

Axiom 2 (Relevance)

If $f(\cdot \mid t) = g(\cdot \mid t)$ for all $t \in S$, then $f \sim_S g$.

- If two lotteries agree on **every state inside** S, they must be indifferent given S
- States **outside** S are irrelevant when choosing conditional on S
- Very natural: you should not care about outcomes in states you know cannot occur

Axiom 3 – Monotonicity

Axiom 3 (Monotonicity)

If $f \succ_S h$ and $0 \leq \beta < \alpha \leq 1$, then:

$$\alpha f + (1 - \alpha)h \succ_S \beta f + (1 - \beta)h$$

- Higher probability of the **better** lottery is always preferred
- Formalizes the intuition that "more of a good thing is better" in probabilistic terms
- Works in conjunction with continuity (Axiom 4) to pin down utility scales

Axiom 4 – Continuity

Axiom 4 (Continuity)

If $f \succsim_S g$ and $g \succsim_S h$, then there exists some $\gamma \in [0,1]$ such that:

$$g \sim_S \gamma f + (1-\gamma)h$$

- Any lottery ranked **between** f and h is equivalent to some mixture of f and h
- Ensures no lottery is "infinitely better" than any other
- Together with monotonicity: preferences vary **continuously** with mixing probabilities
- This axiom is what allows us to represent preferences with **real-valued utility numbers**

Axiom 5A – Objective Substitution

Axiom 5A (Objective Substitution)

If $e \succsim_S f$ and $g \succsim_S h$ and $0 \leq \alpha \leq 1$, then:

$$\alpha e + (1-\alpha) g \succsim_S \alpha f + (1-\alpha) h$$

- The mixing involves **objective randomizations** (coin tosses, roulette wheels with known probabilities)
- If you weakly prefer e over f, and weakly prefer g over h, then mixing preserves rankings
- This is the **von Neumann–Morgenstern independence axiom** – also called the **independence axiom** in most modern texts

Axiom 5B – Strict Objective Substitution

Axiom 5B (Strict Objective Substitution)

If $e \succ_S f$ and $g \succsim_S h$ and $0 < \alpha \leq 1$, then:

$$\alpha e + (1-\alpha)g \succ_S \alpha f + (1-\alpha)h$$

- Strict version: one strict preference is enough to preserve a strict ranking in the mixture
- Together, 5A and 5B generate strong restrictions on preferences **even without the other axioms**
- The Allais paradox is a violation of precisely Axiom 5B

Axiom 6A – Subjective Substitution

Axiom 6A (Subjective Substitution)

If $f \succsim_S g$ and $f \succsim_T g$ and $S \cap T = \emptyset$, then:

$$f \succsim_{S \cup T} g$$

- The "mixing" is over **subjective unknowns** – disjoint subsets of Ω
- Closely related to **Savage's sure-thing principle**: if you prefer f over g in every possible scenario, you must prefer f before learning which scenario occurs
- Why: otherwise you would express a preference you **know** you will want to reverse after observing the event

Axiom 6B – Strict Subjective Substitution

Axiom 6B (Strict Subjective Substitution)

If $f \succ_S g$ and $f \succ_T g$ and $S \cap T = \emptyset$, then:

$$f \succ_{S \cup T} g$$

- Strict version: strictly prefer f in each of two disjoint events $\Rightarrow$ strictly prefer f in their union
- 5A/B handle **objective** mixing; 6A/B handle **subjective** mixing
- We need both because our lottery space L encompasses both objective and subjective uncertainty simultaneously

Axiom 7 – Interest

Axiom 7 (Interest)

For every state $t \in \Omega$, there exist prizes y and z in X such that:

$$[y] \succ_{\{t\}} [z]$$

- A regularity condition: not all prizes are equally valued in every state
- Ensures there is **something at stake** in every state – the utility scale is non-trivial
- Without this, the utility function would be constant and meaningless
- Allows us to define **best** (a_1) and **worst** (a_0) lotteries used in the proof

Axiom 8 – State Neutrality (Optional)

Axiom 8 (State Neutrality)

For any two states $r, t \in \Omega$: if $f(\cdot \mid r) = f(\cdot \mid t)$ and $g(\cdot \mid r) = g(\cdot \mid t)$ and $f \succsim_{\{r\}} g$, then:

$$f \succsim_{\{t\}} g$$

- If two lotteries deliver the **same objective distributions** in states r and t, the preference ranking cannot depend on which state it is
- Fails if the same prize is valued differently in different states (e.g., $\$100$ more valuable when ill than when healthy)
- Adding Axiom 8 yields a **state-independent utility function** $U(x)$

Why Substitution Cannot Be Dropped (I)

Key Idea

Substitution says that a preference should survive inside a lottery.

- Suppose the agent prefers x to y:

$$x \succ y.$$

- Then replacing y by x inside a lottery should make the lottery better.
- The same preference should not disappear just because uncertainty is added.

Why Substitution Cannot Be Dropped (II)

- Suppose substitution fails.
- The agent directly prefers x to y:
$$x \succ y.$$
- But after mixing both prizes with the same prize z:
$$0.5[y] + 0.5[z] \succ [w] \succ 0.5[x] + 0.5[z].$$

What is strange?

The agent likes x more than y directly, but prefers the lottery with y to the one with x.

Why Substitution Cannot Be Dropped (III)

- Consider the following decision problem.

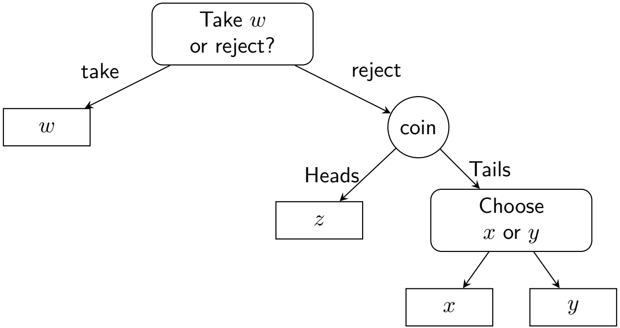

- If the agent rejects w, then:
  - Heads gives z.
  - Tails leads to a direct choice between x and y.

Why Substitution Cannot Be Dropped (IV)

- Before the coin toss, the agent compares plans:

Plan	Lottery
Take w	$[w]$
Reject w; choose x if Tails	$0.5[x] + 0.5[z]$
Reject w; choose y if Tails	$0.5[y] + 0.5[z]$

- Since

$$0.5[y] + 0.5[z] \succ [w] \succ 0.5[x] + 0.5[z],$$

the best plan is:

Reject w and choose y if Tails.

Why Substitution Cannot Be Dropped (V)

- But if Tails occurs, the agent directly chooses between x and y.
- Since $x \succ y$, he chooses x, not y.
- Therefore rejecting w actually gives $0.5[x] + 0.5[z]$, not $0.5[y] + 0.5[z]$.
- But by assumption, $[w] \succ 0.5[x] + 0.5[z]$.

Conclusion

Without substitution, the agent's preferred plan is not credible. He plans to choose y, but later chooses x.

The Expected-Utility Maximization Theorem

Key Objects

- A **conditional-probability function** $p : \Xi \to \Delta(\Omega)$ satisfies:
    - $p(t \mid S) = 0$ if $t \notin S$
    - $\sum_{r \in S} p(r \mid S) = 1$ for every event S
    - Write $p(R \mid S) = \sum_{r \in R} p(r \mid S)$ for $R \subseteq \Omega$

- A **utility function** is any $u : X \times \Omega \to \mathbb{R}$
- It is **state independent** if $u(x, t) = U(x)$ for some $U : X \to \mathbb{R}$

- Given p and u, the **expected utility** of lottery f given event S:

$$E_p(u(f) \mid S) = \sum_{t \in S} p(t \mid S) \sum_{x \in X} u(x, t)\, f(x \mid t) \tag{2}$$

Theorem 1 – Statement

Theorem 1

Axioms 1AB, 2, 3, 4, 5AB, 6AB, 7 are jointly satisfied **if and only if** $\exists$ utility function $u : X \times \Omega \to \mathbb{R}$ and conditional-probability function $p : \Xi \to \Delta(\Omega)$ satisfying conditions (3), (4), (5) below. Furthermore, Axiom 8 is also satisfied iff (3)–(5) can be satisfied with a **state-independent** u.

$$\max_{x \in X} u(x,t) = 1 \quad \text{and} \quad \min_{x \in X} u(x,t) = 0, \quad \forall t \in \Omega \tag{3}$$

$$p(R \mid T) = p(R \mid S) \cdot p(S \mid T), \quad \forall R \subseteq S \subseteq T \subseteq \Omega,\ S \neq \emptyset \tag{4}$$

$$f \succsim_S g \iff E_p(u(f) \mid S) \geq E_p(u(g) \mid S), \quad \forall f, g \in L,\ \forall S \in \Xi \tag{5}$$

Theorem 1.1 – Interpreting the Conditions

- **Condition** (3) – **Normalization:**
  - Utility is scaled to $[0, 1]$ in every state: best prize $= 1$, worst prize $= 0$
  - No loss of generality – always achievable by rescaling
- **Condition** (4) – **Bayes's Formula:**
  - Conditional probabilities are internally consistent
  - This is exactly the **chain rule** of probability
- **Condition** (5) – **Expected-Utility Representation:**
  - Decision-maker always prefers the lottery with **higher expected utility**
  - Once u and p are assessed, choices are predicted in **any** situation
  - With X and Ω finite: infinitely many lotteries in L characterized by **finitely many numbers**

Interpreting the EU Formula – Accounting Details (I)

- The expected utility formula (2) can be rewritten to reveal its probability content:

$$E_p(u(f) \mid S) = \sum_{t \in S} p(t \mid S) \sum_{x \in X} u(x,t)\, f(x \mid t) = \sum_{x \in X} u(x,t) \underbrace{\left[\sum_{t \in S} \overbrace{p(t \mid S)}^{\Pr(t|S)} \; \overbrace{f(x \mid t)}^{\Pr(x|t,f)} \right]}_{\Pr(x|S,f)}$$

- $\Pr(x \mid S, f) = \sum_{t \in S} p(t \mid S)\, f(x \mid t)$ is the **total (unconditional) probability** of prize x under act f, given event S
- This is just the **law of total probability**: average the objective conditional probabilities over states, using subjective state weights

Interpreting the EU Formula – Accounting Details (II)

Note

Intuitively: EU collapses the two-stage lottery into a single-stage one. The agent's subjective beliefs $p(\cdot \mid S)$ and the objective probabilities $f(\cdot \mid t)$ are combined via the law of total probability – and utility is taken as an expectation over the resulting distribution on prizes.

Proof – Intuitive Sketch (I)

- **The Big Idea:** construct a **universal ruler** – two reference acts:
  - a_1: assigns the **degenerate lottery** $[y_t]$ in every state t, where y_t is a best prize in state t
  - a_0: assigns the **degenerate lottery** $[z_t]$ in every state t, where z_t is a worst prize in state t
  - Both are acts in L – the same space as f and g – just very simple benchmark acts
  - Every act can be evaluated relative to these two endpoints
  - Both **utilities** and **probabilities** are measured as positions on the same $[0, 1]$ scale

- **1 – Build the ruler:**
  - a_1 and a_0 exist because X is finite and preferences are complete and transitive
  - The best/worst prize may differ across states – a_1 and a_0 are single acts, but they are state-contingent internally
  - Mixtures of the form $\beta a_1 + (1 - \beta)a_0$ involve only one **objective probability dial** $\beta \in [0, 1]$
  - At this stage, we are not yet using subjective probabilities; we are only building a measurement scale

Proof – Intuitive Sketch (IIA)

- **2 – Elicit utilities:**
  - For each prize x in state t, find β such that:

$$[x] \sim_{\{t\}} \beta a_1 + (1-\beta)a_0$$

  - The right-hand side is a simple benchmark gamble: best prize with objective probability β, worst prize with objective probability $1-\beta$
  - That β **is** $u(x,t)$ – the utility of prize x in state t, measured on the best–worst ruler
  - Continuity guarantees that such a β exists
  - The normalization (3) holds by construction:
    - best prize $\Rightarrow u(y_t, t) = 1$
    - worst prize $\Rightarrow u(z_t, t) = 0$

Proof – Intuitive Sketch (IIB)

- **3 – Elicit probabilities:**
  - For each state t and event S, find γ such that:

$$b_{\{t\}} \sim_S \gamma a_1 + (1-\gamma) a_0$$

  - The left-hand side is a bet on state t
  - The right-hand side is the same best–worst ruler
  - That γ **is** $p(t \mid S)$ – the subjective probability of state t conditional on event S
  - We do not observe this probability directly; we elicit it from the agent's indifference between:
    - betting on state t
    - receiving the best prize with objective probability γ

Proof – Intuitive Sketch (IIB)

- For example, suppose there are three states: $\Omega = \{t_1, t_2, t_3\}$.
- Then the bet on state t_2 is:

$$b_{\{t_2\}} = \begin{cases} a_0, & \text{in state } t_1, \\ a_1, & \text{in state } t_2, \\ a_0, & \text{in state } t_3. \end{cases}$$

- So $b_{\{t_2\}}$ is an act that pays the **best-prize act** only if state t_2 occurs, and pays the **worst-prize act** otherwise.
- To elicit beliefs, compare this subjective bet with an objective best–worst lottery:

$$b_{\{t\}} \sim_S \gamma a_1 + (1-\gamma) a_0.$$

- The left side depends on the agent's belief about t given S; the right side uses the known objective probability γ. Thus the indifference value γ reveals the conditional belief:

$$p(t \mid S) = \gamma.$$

Proof – Intuitive Sketch (III)

- **4 – Every act reduces to the ruler**:
  - **Decompose:** any act f can be broken into elementary prize-state components $c_{x,t}$
  - **Substitute utilities:** each prize x in state t is equivalent to $u(x,t)$ units of the best–worst ruler
  - **Substitute probabilities:** each state t inside event S is weighted according to $p(t \mid S)$
  - **Collect terms:** the act f becomes equivalent to a binary lottery between a_1 and a_0, with weight determined by $E_p(u(f) \mid S)$
  - Therefore, comparing two acts reduces to comparing their conditional expected utilities:

$$f \succsim_S g \quad \Longleftrightarrow \quad E_p(u(f) \mid S) \geq E_p(u(g) \mid S)$$

Note

Axioms 5A/B and 6A/B are the substitution engines of the proof. They say that indifferent pieces can be swapped inside larger acts without changing the overall preference.

Proof – Intuitive Sketch (IV)

- **5 – Bayes's rule falls out for free** (proving (4)):

- Take three events with $R \subseteq S \subseteq T$.
- Evaluate the same bet b_R in two ways, conditional on T

- **Direct valuation:** betting on R given T is equivalent to a best–worst lottery with weight

$$p(R \mid T)$$

- **Two-step valuation:** since $R \subseteq S$, the event R can occur only through S
  - first, S occurs within T
  - then, R occurs within S

The corresponding belief weights are:

$$p(S \mid T) \quad \text{and} \quad p(R \mid S)$$

Proof – Intuitive Sketch (V)

- Therefore, the two-step evaluation gives the weight

$$p(R \mid S)p(S \mid T)$$

- But both procedures evaluate the same bet b_R conditional on T

- Hence consistency requires:

$$p(R \mid T) = p(R \mid S)p(S \mid T)$$

- So Bayes's rule is not imposed from outside:
  - it follows from preference consistency
  - the same bet cannot receive two different best–worst ruler positions
  - the **chain rule is a theorem**, not an assumption

Proof – Intuitive Sketch (VI)

- **The unifying insight:**
  - The proof does not assume expected utility or subjective probability
  - It **constructs** both from preferences
  - The best and worst acts a_1 and a_0 provide a common ruler
  - Utilities measure prizes on that ruler
  - Probabilities measure beliefs on that ruler, through betting indifferences
  - Expected utility appears because every act can be reduced to its position on the same ruler
  - Bayes's rule appears because the same bet must have the same position whether evaluated directly or in two consistent steps

Note

The entire proof is a measurement argument. Choices reveal utilities and beliefs because the axioms make all comparisons consistent with one common best–worst scale.

Equivalent Representations

Motivation – Non-Uniqueness

- When condition (3) is dropped, multiple pairs (u, p) may represent the same preferences
- Equivalent representations are **observationally indistinguishable** in terms of decision-theoretic properties
- Any theory requiring us to distinguish between them is suspect

Definition: Representation

(v, q) **represent** the preference ordering $\succsim_S$ if for every $f, g \in L$:

$$E_q(v(f) \mid S) \geq E_q(v(g) \mid S) \iff f \succsim_S g$$

Theorem 2 – General Equivalence

Theorem 2

Suppose preferences satisfy Axioms 1–7, and (u, p) satisfies (3)–(5). Then (v, q) represents $\succsim_S$ **if and only if** $\exists\, A > 0$ and $B : S \to \mathbb{R}$ such that:

$$q(t \mid S)\, v(x, t) = A\, p(t \mid S)\, u(x, t) + B(t), \quad \forall t \in S,\ \forall x \in X \tag{6}$$

- $q(\cdot \mid S)$ can be almost anything – probabilities and utilities are only **jointly identified**
- Without state neutrality, we cannot separately pin down beliefs and tastes

Proof of Theorem 2 – Direction ($\Leftarrow$)

- **Assume** (6) **holds. Show** (v, q) **represents the same preferences.**
  - For any lottery f, compute $E_q(v(f) \mid S)$:

$$E_q(v(f) \mid S) = \sum_{t \in S} \sum_{x \in X} f(x \mid t)\, q(t \mid S)\, v(x, t)$$

  - Substitute (6): $q(t \mid S) v(x, t) = A p(t \mid S) u(x, t) + B(t)$

$$= \sum_{t \in S} \sum_{x \in X} f(x \mid t) \big[A\, p(t \mid S)\, u(x, t) + B(t) \big]$$

$$= A\, E_p(u(f) \mid S) + \sum_{t \in S} B(t) \underbrace{\sum_{x \in X} f(x \mid t)}_{=1} = A\, E_p(u(f) \mid S) + \sum_{t \in S} B(t)$$

Proof of Theorem 2 – Direction ($\Rightarrow$) (I)

- **Assume (v, q) represents the same preferences. Construct A and $B(t)$.**
  - Define:

$$\lambda = \frac{E_q(v(c_{x,t}) \mid S) - E_q(v(a_0) \mid S)}{E_q(v(a_1) \mid S) - E_q(v(a_0) \mid S)}$$

  - By linearity of expectations: $E_q(v(\lambda a_1 + (1-\lambda)a_0) \mid S) = E_q(v(c_{x,t}) \mid S)$
  - So $c_{x,t} \sim_S \lambda a_1 + (1-\lambda)a_0$
  - But from the proof of Theorem 1, we constructed u and p so that:

$$c_{x,t} \sim_S p(t \mid S)\, u(x,t)\, a_1 + (1 - p(t \mid S)\, u(x,t))\, a_0$$

  - By monotonicity, only one mixing probability can be indifferent to $c_{x,t}$, so: $\lambda = p(t \mid S)\, u(x,t)$

Proof of Theorem 2 – Direction ($\Rightarrow$) (II)

- Note that $c_{x,t}$ differs from a_0 only in state t, where it gives prize x instead of the worst prize:

$$E_q(v(c_{x,t}) \mid S) - E_q(v(a_0) \mid S) = q(t \mid S)\Big(v(x,t) - \min_{z \in X} v(z,t)\Big)$$

- Going back to the definition of λ:

$$p(t \mid S)\, u(x,t) = \frac{q(t \mid S)(v(x,t) - \min_z v(z,t))}{E_q(v(a_1) \mid S) - E_q(v(a_0) \mid S)}$$

- Set $A = E_q(v(a_1) \mid S) - E_q(v(a_0) \mid S) > 0$ (positive since $a_1 \succ_S a_0$)
- Set $B(t) = q(t \mid S) \min_z v(z,t)$
- Then $q(t \mid S)v(x,t) = Ap(t \mid S)u(x,t) + B(t)$, i.e., (6) holds ■

Theorem 3 – Uniqueness with State Independence

Theorem 3

Suppose Axioms 1–8 hold. Let (u, p) be the state-independent representation from Theorem 1. If (v, q) is **any** state-independent representation of $\succsim_S$, then:

$$q(t \mid S) = p(t \mid S), \quad \forall t \in S$$

and $\exists\, A > 0$, $C \in \mathbb{R}$ such that $v(x) = Au(x) + C$, $\forall x \in X$.

- State neutrality **uniquely pins down** the probability distribution
- Utility is unique **up to a positive affine transformation** – the standard vNM cardinality result

Proof of Theorem 3

- Set $A = E_q(v(a_1) \mid S) - E_q(v(a_0) \mid S)$ and $C = \min_{z \in X} v(z)$
- From the proof of Theorem 2 with state-independent u and v:

$$A\, p(t \mid S)\, u(x) + q(t \mid S)\, C = q(t \mid S)\, v(x), \quad \forall x \in X,\ \forall t \in S$$

- Sum over all $t \in S$ and use $\sum_{t \in S} p(t \mid S) = \sum_{t \in S} q(t \mid S) = 1$:

$$A\, u(x) + C = v(x)$$

- Substitute $v(x) = Au(x) + C$ back into the equation for a single t, setting x to the best prize (so $u(x) = 1$):

$$A\, p(t \mid S) + q(t \mid S)\, C = A\, q(t \mid S) + q(t \mid S)\, C$$

- Since $A > 0$: $p(t \mid S) = q(t \mid S)$ ■

Bayesian Conditional-Probability Systems

Definition

Definition: Bayesian Conditional-Probability System

$p \in \Delta^*(\Omega)$ if p is a conditional-probability function on Ω satisfying condition (4) (Bayes's formula):

$$p(R \mid T) = p(R \mid S)\, p(S \mid T), \quad \forall R \subseteq S \subseteq T \subseteq \Omega,\ S \neq \emptyset$$

- Let $\Delta^0(Z) = \{q \in \Delta(Z) \mid q(z) > 0,\ \forall z \in Z\}$: full-support distributions
- Any $\hat{p} \in \Delta^0(\Omega)$ generates a Bayesian conditional-probability system via:
  - $p(t \mid S) = \dfrac{\hat{p}(t)}{\sum_{r \in S} \hat{p}(r)}$ if $t \in S$, $\quad p(t \mid S) = 0$ if $t \notin S$
- But $\Delta^*(\Omega)$ is strictly larger than what full-support priors generate

Theorem 4 – Characterization of $\Delta^*(\Omega)$

Important

Theorem 4 (Myerson, 1986b)
$p \in \Delta^*(\Omega)$ **if and only if** there exists a sequence $\{\hat{p}^k\}_{k=1}^{\infty}$ in $\Delta^0(\Omega)$ such that for every nonempty $S \subseteq \Omega$ and every $t \in \Omega$:

$$p(t \mid S) = \lim_{k\to\infty} \frac{\hat{p}^k(t)}{\sum_{r\in S} \hat{p}^k(r)} \text{ if } t \in S, \qquad p(t \mid S) = 0 \text{ if } t \notin S$$

- Any Bayesian conditional-probability system is a **limit** of standard Bayesian conditioning from full-support priors
- This justifies assigning beliefs after **zero-probability events** (off-equilibrium paths)
- Essential foundation for sequential equilibrium and perfect Bayesian equilibrium refinements

A Simple Example of $\Delta^*(\Omega)$ (I)

Setup

$\Omega = \{t_1, t_2, t_3\}$, prior $\hat{p} = (1/2,\ 1/2,\ 0)$ – state t_3 has zero probability

- **Approximating sequence:** for $k = 1, 2, 3, \ldots$

$$\hat{p}^k = \left(\frac{1}{2} - \frac{1}{k},\quad \frac{1}{2} - \frac{1}{k},\quad \frac{2}{k}\right) \in \Delta^0(\Omega)$$

- **Conditioning on the zero-probability event** $S = \{t_3\}$:

$$p^k(t_3 \mid \{t_3\}) = \frac{\hat{p}^k(t_3)}{\hat{p}^k(t_3)} = \frac{2/k}{2/k} = 1 \quad \forall k$$

A Simple Example of $\Delta^*(\Omega)$ (II)

- Standard Bayes fails: $\hat{p}(t_3) = 0$ so $\dfrac{\hat{p}(t_3)}{\hat{p}(t_3)}$ is $\frac{0}{0}$
- The limit recovers the intuitive answer: **if told t_3 occurred, believe it**
- This is precisely what $\Delta^*(\Omega)$ allows – coherent beliefs after zero-probability events

Limitations of the Bayesian Model

Descriptive vs. Prescriptive Validity

Descriptive	Prescriptive
Predict what people *actually* do	Guide what people *should* do
Tested by experimental and empirical data	Tested by whether a rational agent would regret deviating
Systematic violations documented in experiments	Theorem 1 provides a proof of prescriptive validity

- If a model is prescriptively valid, deviations are **mistakes**
- With sufficient time to learn, agents make fewer mistakes
- So EU maximization should also be **predictively accurate** in many settings

The Allais Paradox – Setup

- $X = \{\$12M, \$1M, \$0\}$; define four lotteries:
    - $f_1 = .10[\$12M] + .90[\$0]$
    - $f_2 = .11[\$1M] + .89[\$0]$
    - $f_3 = [\$1M]$
    - $f_4 = .10[\$12M] + .89[\$1M] + .01[\$0]$

The Allais Paradox – Setup

- $X = \{\$12M, \$1M, \$0\}$; define four lotteries:
  - $f_1 = .10[\$12M] + .90[\$0]$
  - $f_2 = .11[\$1M] + .89[\$0]$
  - $f_3 = [\$1M]$
  - $f_4 = .10[\$12M] + .89[\$1M] + .01[\$0]$

- Common preferences: $f_1 \succ f_2$ and $f_3 \succ f_4$
- Reasoning for $f_1 \succ f_2$: $\$12M$ is so much better that the slightly higher winning probability in f_2 is not worth the lower prize
- Reasoning for $f_3 \succ f_4$: sure $\$1M$ beats accepting a 1% chance of $\$0$ in exchange for a 10% chance at $\$12M$

The Allais Paradox – Proof of Inconsistency

- Compute the compound lotteries:

$$0.5f_1 + 0.5f_3 = 0.05[\$12M] + 0.5[\$1M] + 0.45[\$0]$$

$$0.5f_2 + 0.5f_4 = 0.5(0.11[\$1M] + 0.89[\$0]) + 0.5(0.10[\$12M] + 0.89[\$1M] + 0.01[\$0])$$

$$= 0.05[\$12M] + 0.5[\$1M] + 0.45[\$0]$$

- The two compound lotteries are **identical**: $0.5f_1 + 0.5f_3 = 0.5f_2 + 0.5f_4$
- But by Axiom 5B (strict objective substitution):
  - $f_1 \succ f_2$ and $f_3 \succ f_4 \Rightarrow 0.5f_1 + 0.5f_3 \succ 0.5f_2 + 0.5f_4$
  - This contradicts the identical lotteries $\Rightarrow$ **no utility function can rationalize these preferences** ■

The Ellsberg Paradox – Urn Setup

- An urn contains **90 balls** of three colors:
  - **30 red** balls – known exactly
  - **60 black or yellow** balls – unknown proportion

- Four bets, each paying $\$100$ on the relevant draw, $\$0$ otherwise:

Bet	Red	Black	Yellow
f_1	$100	$0	$0
f_2	$0	$100	$0
f_3	$100	$0	$100
f_4	$0	$100	$100

- Note: f_3 and f_4 are complements – exactly one of them pays $\$100$ in every draw

The Ellsberg Paradox – Typical Preferences

- Most people who face this problem choose:

$$f_1 \succ f_2 \qquad \text{and} \qquad f_4 \succ f_3$$

- **Intuition behind each choice:**
  - $f_1 \succ f_2$: red is known to be $\frac{1}{3}$ of the urn; black is unknown – prefer the known risk
  - $f_4 \succ f_3$: black+yellow covers 60 balls for sure; red+yellow has unknown yellow component – again prefer the less ambiguous bet

> **Note**
>
> Both choices reflect **ambiguity aversion** – a preference for known probabilities over unknown ones, regardless of which side of the bet you are on

The Ellsberg Paradox – Violation of SEU (I)

- **Suppose** a subjective probability p over $\{R, B, Y\}$ rationalizes $f_1 \succ f_2$:

$$f_1 \succ f_2 \implies E_p(u(f_1)) > E_p(u(f_2))$$

$$\implies u(\$100)\, p(R) > u(\$100)\, p(B)$$

$$\implies p(R) > p(B)$$

- Since there are 30 red balls out of 90: $p(R) = \frac{1}{3}$ is objectively pinned down
- Therefore: $f_1 \succ f_2 \implies p(B) < \frac{1}{3}$

The Ellsberg Paradox – Violation of SEU (II)

- **Now suppose** the same p rationalizes $f_4 \succ f_3$:

$$f_4 \succ f_3 \implies E_p(u(f_4)) > E_p(u(f_3))$$
$$\implies u(\$100)\,[p(B) + p(Y)] > u(\$100)\,[p(R) + p(Y)]$$
$$\implies p(B) > p(R) = \tfrac{1}{3}$$

- But from the previous slide: $p(B) < \frac{1}{3}$

Contradiction

$$f_1 \succ f_2 \implies p(B) < \tfrac{1}{3} \qquad \text{and} \qquad f_4 \succ f_3 \implies p(B) > \tfrac{1}{3}$$

No single probability distribution over $\{R, B, Y\}$ can rationalize both preferences simultaneously ■

The Ellsberg Paradox – Violation of Axiom 5B

- There is a second, more direct way to see the violation – through **objective substitution**
- Observe that f_1 and f_2 mix to exactly f_3 and f_4:

$$\frac{1}{2}f_1 + \frac{1}{2}f_4 = \frac{1}{2}f_2 + \frac{1}{2}f_3$$

- Both sides pay $\$100$ with probability $\frac{1}{2}$ regardless of which ball is drawn – verify:

	Red	**Black**	**Yellow**
$\frac{1}{2}f_1 + \frac{1}{2}f_4$	$\frac{1}{2}$(\$100) + $\frac{1}{2}$(\$100)	$\frac{1}{2}$(\$0) + $\frac{1}{2}$(\$100)	$\frac{1}{2}$(\$0) + $\frac{1}{2}$(\$100)
$\frac{1}{2}f_2 + \frac{1}{2}f_3$	$\frac{1}{2}$(\$0) + $\frac{1}{2}$(\$100)	$\frac{1}{2}$(\$100) + $\frac{1}{2}$(\$0)	$\frac{1}{2}$(\$100) + $\frac{1}{2}$(\$0)

The Ellsberg Paradox – Violation of Axiom 5B (Cont.)

- Both mixtures yield **identical objective lotteries** in every state – yet:

$$f_1 \succ f_2 \quad \text{and} \quad f_4 \succ f_3$$

- By Axiom 5B applied twice:

$$\tfrac{1}{2}f_1 + \tfrac{1}{2}f_4 \succ \tfrac{1}{2}f_2 + \tfrac{1}{2}f_4 \succ \tfrac{1}{2}f_2 + \tfrac{1}{2}f_3$$

- But we just showed $\tfrac{1}{2}f_1 + \tfrac{1}{2}f_4 = \tfrac{1}{2}f_2 + \tfrac{1}{2}f_3$, so this says:

$$\tfrac{1}{2}f_1 + \tfrac{1}{2}f_4 \succ \tfrac{1}{2}f_1 + \tfrac{1}{2}f_4$$

Contradiction

An act cannot be strictly preferred to itself. The Ellsberg preferences violate Axiom 5B directly – **before** any probability representation is even invoked

Domination

Motivation – Probability-Free Statements

- In games, assessing subjective probabilities is especially hard
- The "state" may include other players' decisions, which depend on your decisions – a **circular structure**
- But sometimes we can say an option is **never optimal** regardless of beliefs

- Let $u : X \times \Omega \to \mathbb{R}$; the agent chooses $x \in X$
- With beliefs $p \in \Delta(\Omega)$, option $y \in X$ is **optimal** iff:

$$\sum_{t \in \Omega} p(t)\, u(y,t) \geq \sum_{t \in \Omega} p(t)\, u(x,t), \quad \forall x \in X \tag{7}$$

Theorem 5 – Convexity

Important

Theorem 5
Given $u : X \times \Omega \to \mathbb{R}$ and $y \in X$, the set of all $p \in \Delta(\Omega)$ such that y is optimal is **convex**.

- **Proof:** Suppose y is optimal under beliefs p and under beliefs q; let $r = \lambda p + (1 - \lambda)q$, $\lambda \in [0, 1]$:

$$\sum_t r(t)\, u(y,t) = \lambda \sum_t p(t)\, u(y,t) + (1-\lambda) \sum_t q(t)\, u(y,t)$$

$$\geq \lambda \sum_t p(t)\, u(x,t) + (1-\lambda) \sum_t q(t)\, u(x,t) = \sum_t r(t)\, u(x,t)$$

- So y is optimal under $r = \lambda p + (1 - \lambda)q$ ■

Domination – Definitions

- A **randomized strategy** is any $\sigma \in \Delta(X)$, where $\sigma(x)$ = probability of choosing x

Definition: Strong Domination

$y \in X$ is **strongly dominated** by $\sigma \in \Delta(X)$ iff:

$$\sum_{x \in X} \sigma(x)\, u(x,t) > u(y,t), \quad \forall t \in \Omega \tag{8}$$

- σ gives strictly higher expected payoff than y in **every** state

Weak Domination – Definition

Definition: Weak Domination

$y \in X$ is **weakly dominated** by $\sigma \in \Delta(X)$ iff:

$$\sum_{x \in X} \sigma(x)\, u(x,t) \geq u(y,t) \quad \forall t \in \Omega$$

and there exists at least one state s such that:

$$\sum_{x \in X} \sigma(x)\, u(x,s) > u(y,s)$$

- σ is never worse than y in any state, and strictly better in at least one state
- Weakly dominated options can be optimal – but only when some states receive probability zero

Worked Example

- **Setup:** $X = \{\alpha, \beta, \gamma\}$, $\Omega = \{\theta_1, \theta_2\}$

Decision	θ_1	θ_2
α	8	1
β	5	3
γ	4	7

- Let $p(\theta_1) = p$, so $p(\theta_2) = 1 - p$
- Each decision is optimal for some **convex interval** of beliefs (Theorem 5)

Worked Example – Optimality Regions

- α **optimal** when both hold:
  - $8p + 1(1-p) \geq 5p + 3(1-p) \Rightarrow 3p \geq 2(1-p) \Rightarrow p \geq 2/5$
  - $8p + 1(1-p) \geq 4p + 7(1-p) \Rightarrow 4p \geq 6(1-p) \Rightarrow p \geq 3/5$
  - $\Rightarrow \alpha$ **optimal when** $p \geq 0.6$
- γ **optimal** when both hold:
  - $4p + 7(1-p) \geq 8p + 1(1-p) \Rightarrow p \leq 3/5$
  - $4p + 7(1-p) \geq 5p + 3(1-p) \Rightarrow p \leq 1$ (always)
  - $\Rightarrow \gamma$ **optimal when** $p \leq 0.6$
- β **optimal** requires $5p + 3(1-p) \geq 8p + 1(1-p)$ $(p \leq 2/5)$ AND $5p + 3(1-p) \geq 4p + 7(1-p)$ $(p \geq 1)$ – **no** p **satisfies both** $\Rightarrow \beta$ **is strongly dominated**

Worked Example – β Dominated by a Randomized Strategy

- Consider $\sigma = 0.5[\alpha] + 0.5[\gamma]$:

State	σ expected payoff	β payoff	Strictly better?
θ_1	$0.5 \times 8 + 0.5 \times 4 = 6$	5	Yes: $6 > 5$
θ_2	$0.5 \times 1 + 0.5 \times 7 = 4$	3	Yes: $4 > 3$

- σ beats β in **every state** – regardless of beliefs
- So β satisfies condition (8): it is strongly dominated by σ
- The dominating strategy σ need not itself be optimal – it only needs to beat β

Weak Domination – Example

Setup: $X = \{\alpha, \beta\}$, $\Omega = \{\theta_1, \theta_2\}$

Decision	θ_1	θ_2
α	5	3
β	5	1

- β is **weakly dominated** by α ($=$ the pure strategy $[\alpha]$):
  - In θ_1: α and β tie $(5 = 5)$
  - In θ_2: α strictly beats β $(3 > 1)$
- β is optimal only if $p(\theta_1) = 1$ (probability one on the tie state)
- Any positive probability on θ_2 makes α strictly preferred
- By Theorem 7: no $p \in \Delta^0(\Omega)$ makes β optimal $\iff$ β is weakly dominated

Summary of Main Results

Result	Content
Thm 1	Axioms 1–7 $\iff$ EU maximization; add 8 for state-independent utility
Thm 2	(v, q) represents same preferences as (u, p) iff $q(t\|S)v(x, t) = Ap(t\|S)u(x, t) + B(t)$
Thm 3	With state neutrality: probability unique; utility unique up to positive affine transform
Thm 4	Every Bayesian CPS is a limit of standard Bayesian updating from full-support priors
Thm 5	Set of beliefs making y optimal is convex

Required Reading

- Myerson, R.B. (1991). *Game Theory: Analysis of Conflict*. **Chapter 1: Sections 1.1–1.9** (pp. 1–36)

- Papers:
  - Anscombe, F.J., and Aumann, R.J. (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34(1), 199–205.
  - Samuelson, L. (2016). Game Theory in Economics and Beyond. *Journal of Economic Perspectives*, 30(4), 107–130.