International School of Economics at TSU
Microeconomics IV (Game Theory)
Lasha Chochua

Problem Set 1 – Solutions

Problem 1 – Lotteries and the Formal Setting

Let $\Omega = \{t_1, t_2\}$ and $X = \{\$0, \$4, \$9\}$. Consider the following two acts:

$$f = \begin{bmatrix} & \$0 & \$4 & \$9 \\ t_1 & 0.5 & 0.5 & 0 \\ t_2 & 0 & 0.4 & 0.6 \end{bmatrix} \qquad g = \begin{bmatrix} & \$0 & \$4 & \$9 \\ t_1 & 0.2 & 0.6 & 0.2 \\ t_2 & 0.1 & 0.3 & 0.6 \end{bmatrix}$$

(a) Verify that f and g are valid lotteries, i.e. that both belong to $L = \{h : \Omega \to \Delta(X)\}$.

(b) Write out the compound lottery $k = 0.4f + 0.6g$ explicitly as a matrix. Verify that $k \in L$.

(c) Write out the degenerate lottery [\$4] as a matrix. Does $[\$4] \in L$? Is [\$4] a special case of a compound lottery? Explain.

(d) Let $S = \{t_1\}$. Axiom 2 (Relevance) states that if two acts agree on every state in S, they must be indifferent given S. Construct an act $\tilde{f}$ that differs from f only outside S, and state precisely what Axiom 2 implies about f and $\tilde{f}$.

(a) Solution. A lottery $f \in L$ requires that for every state $t \in \Omega$, the row $f(\cdot \mid t)$ is a valid probability distribution over X: all entries non-negative, summing to 1.

For f: row t_1 has entries $0.5, 0.5, 0 \geq 0$ summing to 1; row t_2 has entries $0, 0.4, 0.6 \geq 0$ summing to 1. So $f \in L$.

For g: row t_1 has entries $0.2, 0.6, 0.2 \geq 0$ summing to 1; row t_2 has entries $0.1, 0.3, 0.6 \geq 0$ summing to 1. So $g \in L$.

(b) Solution. The compound lottery is defined entry-by-entry: $k(x \mid t) = 0.4\, f(x \mid t) + 0.6\, g(x \mid t)$.

$$k = \begin{bmatrix} & \$0 & \$4 & \$9 \\ t_1 & 0.4(0.5) + 0.6(0.2) & 0.4(0.5) + 0.6(0.6) & 0.4(0) + 0.6(0.2) \\ t_2 & 0.4(0) + 0.6(0.1) & 0.4(0.4) + 0.6(0.3) & 0.4(0.6) + 0.6(0.6) \end{bmatrix} = \begin{bmatrix} & \$0 & \$4 & \$9 \\ t_1 & 0.32 & 0.56 & 0.12 \\ t_2 & 0.06 & 0.34 & 0.60 \end{bmatrix}$$

Row t_1: $0.32 + 0.56 + 0.12 = 1$. Row t_2: $0.06 + 0.34 + 0.60 = 1$. All entries non-negative. So $k \in L$.

(c) Solution. The degenerate lottery [\$4] gives prize \$4 with certainty in every state:

$$[\$4] = \begin{bmatrix} & \$0 & \$4 & \$9 \\ t_1 & 0 & 1 & 0 \\ t_2 & 0 & 1 & 0 \end{bmatrix}$$

Yes, $[\$4] \in L$: each row sums to 1 and all entries are non-negative. It is a special case of a compound lottery. For any $\alpha \in [0,1]$ and any $h \in L$, the compound lottery $\alpha\,[\$4]+(1-\alpha)\,[\$4] = [\$4]$, so it is trivially expressible as a mixture. More substantively, a degenerate lottery is the limiting case in which the objective randomization places all weight on a single prize – i.e., $f(x \mid t) = 1$ for exactly one x and zero otherwise, in every state.

(d) Solution. Define:

$$\tilde{f} = \begin{bmatrix} & \$0 & \$4 & \$9 \\ t_1 & 0.5 & 0.5 & 0 \\ t_2 & 1 & 0 & 0 \end{bmatrix}$$

This agrees with f on state t_1 (same row) but differs on t_2. Since $f(\cdot \mid t_1) = \tilde{f}(\cdot \mid t_1)$, the two acts agree on every state in $S = \{t_1\}$. Axiom 2 (Relevance) then implies $f \sim_{\{t_1\}} \tilde{f}$: conditional on the state being t_1, the agent is indifferent between f and $\tilde{f}$, even though they differ drastically on t_2. When we condition on $S = \{t_1\}$, state t_2 is known not to occur and is therefore irrelevant to the ranking.

Problem 2 – The Axioms

(a) Suppose $f \succ_S g$ and $g \sim_S h$. Using only the definitions of $\succ_S$ and $\sim_S$ in terms of $\succsim_S$, and Axioms 1A and 1B, prove that $f \succ_S h$.

(b) Suppose $f \succ_S h$ and $\alpha > \beta$, both in $[0,1]$. Axiom 3 states that $\alpha f + (1-\alpha)h \succ_S \beta f + (1-\beta)h$. Give an intuitive explanation of why a rational agent must satisfy this axiom. What would it mean for a decision-maker's behavior if Axiom 3 were violated?

(c) Explain the role of Axiom 4 in allowing preferences to be represented by real-valued utility numbers. Why is it called a continuity axiom? What type of preferences would violate it? Give a concrete example.

(d) State Axiom 6A precisely. Then explain: if $f \succsim_S g$ and $f \succsim_T g$ for disjoint events S and T, why would it be irrational to have $g \succ_{S \cup T} f$? Connect your answer to the notion of a self-consistent theory.

(a) Solution. We want to show $f \succ_S h$, i.e. $f \succsim_S h$ and $\neg(h \succsim_S f)$.

Recall: $f \succ_S g$ means $f \succsim_S g$ and $\neg(g \succsim_S f)$; and $g \sim_S h$ means $g \succsim_S h$ and $h \succsim_S g$.

Step 1: Show $f \succsim_S h$.

From $f \succ_S g$: $f \succsim_S g$ (by definition of $\succ_S$).

From $g \sim_S h$: $g \succsim_S h$ (by definition of $\sim_S$).

By Axiom 1B (Transitivity): $f \succsim_S g$ and $g \succsim_S h$ imply $f \succsim_S h$.

Step 2: Show $\neg(h \succsim_S f)$.

Suppose for contradiction that $h \succsim_S f$.

From $g \sim_S h$: $g \succsim_S h$ (by definition of $\sim_S$).

By Axiom 1B: $g \succsim_S h$ and $h \succsim_S f$ imply $g \succsim_S f$.

But $f \succ_S g$ means $\neg(g \succsim_S f)$ – contradiction.

Therefore $\neg(h \succsim_S f)$.

Conclusion: $f \succsim_S h$ and $\neg(h \succsim_S f)$, so $f \succ_S h$. ■

(b) Solution. If $f \succ_S h$, then f is strictly better than h. A mixture $\alpha f + (1-\alpha)h$ can be thought of as a lottery that delivers the better option with probability α and the worse option with probability $1-\alpha$. Monotonicity says that raising α – giving more weight to the better option – must strictly improve the mixture. This is a minimal consistency requirement: it would be incoherent to prefer f over h yet be indifferent (or prefer the reverse) when the probability of f increases.

A violation of Axiom 3 would mean the agent strictly prefers f to h but is indifferent (or prefers) a lottery that gives f with lower probability. This immediately undermines any expected utility representation, since EU is linear in probabilities: $EU(\alpha f + (1-\alpha)h) = \alpha \cdot EU(f) + (1-\alpha) \cdot EU(h)$, which is strictly increasing in α whenever $EU(f) > EU(h)$. A violation of Axiom 3 would therefore rule out an expected-utility representation that is linear in objective probabilities. It would mean that the agent's ranking of mixtures is not consistent with the ranking of the underlying alternatives.

(c) Solution. Axiom 4 states: if $f \succsim_S g$ and $g \succsim_S h$, then there exists $\gamma \in [0,1]$ such that $g \sim_S \gamma f + (1-\gamma)h$. The axiom is called a continuity axiom because it rules out discontinuous jumps in the preference ordering. Specifically, it says that any lottery g ranked strictly between the best (f) and worst (h) options can be matched exactly by some mixture of f and h – so preferences vary continuously with mixing probabilities. This is what allows us to assign to each lottery a real-valued index, interpreted as the probability of the best reference outcome that makes the agent indifferent. Without continuity, such a number might not exist.

A classic violation is **lexicographic preferences**. Suppose the agent ranks lotteries first by their probability of \$100, and breaks ties by their probability of \$50. Consider $f = [\$100]$ (the best), $h = [\$0]$ (the worst), and $g = 0.4[\$100] + 0.6[\$50]$. Then g is ranked strictly below f and strictly above h in the lexicographic order. But any mixture $\gamma f + (1-\gamma)h = \gamma[\$100] + (1-\gamma)[\$0]$ gives \$100 with probability γ and \$0 otherwise, so it ranks above g whenever $\gamma > 0.4$ and below

g whenever $\gamma < 0.4$. At $\gamma = 0.4$ exactly, the mixture gives \$100 with probability 0.4 and \$0 with probability 0.6, which ranks strictly below g in the lexicographic order (same probability of \$100, but no \$50). So no γ achieves $g \sim \gamma f + (1-\gamma)h$, violating Axiom 4.

(d) Solution. Axiom 6A states: for any $f, g \in L$ and disjoint events $S, T \in \Xi$, if $f \succsim_S g$ and $f \succsim_T g$, then $f \succsim_{S \cup T} g$.

If $g \succ_{S\cup T} f$ while $f \succsim_S g$ and $f \succsim_T g$, the agent is saying he prefers g over the full event $S \cup T$, yet also says f is at least as good as g in every sub-scenario – both when the state is in S and when it is in T. But the true state must be in S or in T (since they are disjoint and together cover $S \cup T$). Before learning which sub-event occurred, the agent prefers g; the moment he learns which one, he prefers f. This is a preference he knows in advance he will want to reverse. A theory that predicts such behavior cannot survive being understood by the agents it describes: a rational agent, foreseeing the reversal, would act on $f \succsim g$ immediately. Axiom 6A is therefore not an empirical restriction but a self-consistency requirement – it rules out preferences that the agent himself would recognize as incoherent.

Problem 3 – The Substitution Axioms and the Allais Paradox

Recall the four Allais lotteries (prizes in millions of dollars):

$$f_1 = 0.10[\$12] + 0.90[\$0], \quad f_2 = 0.11[\$1] + 0.89[\$0],$$

$$f_3 = [\$1], \quad f_4 = 0.10[\$12] + 0.89[\$1] + 0.01[\$0].$$

(a) Compute $0.5f_1 + 0.5f_3$ and $0.5f_2 + 0.5f_4$ explicitly as probability distributions over $\{\$0, \$1, \$12\}$. Verify they are identical. Show that the common preference pattern $f_1 \succ f_2$ and $f_3 \succ f_4$ violates Axiom 5B.

(b) Suppose instead a decision-maker has preferences $f_2 \succ f_1$ and $f_3 \succ f_4$. Is this consistent with the substitution axioms? If so, find the conditions on u (normalized with $u(\$0) = 0$ and $u(\$12) = 1$) that rationalize both preferences simultaneously.

(c) Suppose $x \succ y$ but $0.5[y] + 0.5[z] \succ [w] \succ 0.5[x] + 0.5[z]$, in violation of Axiom 5B. An agent first decides whether to take prize w or not; if not, a fair coin is tossed, and Heads gives z while Tails gives a choice between x and y. Show that the three behavioral assumptions – (i) can commit, (ii) cannot commit but foresees the inconstancy, (iii) cannot commit and cannot foresee – lead to three different outcomes.

(d) The Allais paradox is sometimes defended as rational behavior under regret theory. In two to three sentences, explain why Myerson's framework rejects this defense.

(a) Solution.

$0.5f_1 + 0.5f_3$:

$$0.5(0.10[\$12] + 0.90[\$0]) + 0.5[\$1] = 0.05[\$12] + 0.45[\$0] + 0.50[\$1]$$
$$= 0.45[\$0] + 0.50[\$1] + 0.05[\$12].$$

$0.5f_2 + 0.5f_4$:

$$0.5(0.11[\$1] + 0.89[\$0]) + 0.5(0.10[\$12] + 0.89[\$1] + 0.01[\$0])$$
$$= 0.055[\$1] + 0.445[\$0] + 0.05[\$12] + 0.445[\$1] + 0.005[\$0]$$
$$= (0.445 + 0.005)[\$0] + (0.055 + 0.445)[\$1] + 0.05[\$12]$$
$$= 0.45[\$0] + 0.50[\$1] + 0.05[\$12].$$

The two compound lotteries are identical. Now suppose $f_1 \succ f_2$ and $f_3 \succ f_4$. Axiom 5B states: if $e \succ f$ and $g \succsim h$ with $\alpha \in (0, 1]$, then $\alpha e + (1-\alpha)g \succ \alpha f + (1-\alpha)h$. Applying with $\alpha = 0.5$, $e = f_1$, $f = f_2$, $g = f_3$, $h = f_4$: since $f_1 \succ f_2$ and $f_3 \succ f_4$ (so in particular $f_3 \succsim f_4$), Axiom 5B gives $0.5f_1 + 0.5f_3 \succ 0.5f_2 + 0.5f_4$. But these two lotteries are identical, so we get the contradiction that a lottery is strictly preferred to itself. ■

(b) Solution. Let $v = u(\$1\text{M}) \in (0, 1)$, with $u(\$0) = 0$ and $u(\$12\text{M}) = 1$.

$f_2 \succ f_1$ requires $EU(f_2) > EU(f_1)$:

$$0.11v + 0.89(0) > 0.10(1) + 0.90(0) \implies 0.11v > 0.10 \implies v > \tfrac{10}{11}.$$

$f_3 \succ f_4$ requires $EU(f_3) > EU(f_4)$:

$$v > 0.10(1) + 0.89v + 0.01(0) \implies v - 0.89v > 0.10 \implies 0.11v > 0.10 \implies v > \tfrac{10}{11}.$$

Both preferences impose exactly the same condition: $u(\$1\text{M}) > \frac{10}{11} \approx 0.909$. This is fully consistent with expected utility theory. It corresponds to a decision-maker who values \$1M very highly relative to \$12M – i.e., exhibits strong diminishing marginal utility of wealth – and is therefore willing to pay a small probability premium to secure \$1M rather than gamble for \$12M.

(c) Solution. The agent's preferences by assumption: $x \succ y$ and $0.5[y] + 0.5[z] \succ [w] \succ 0.5[x] + 0.5[z]$.

The three strategies available are: take w (strategy 1); reject w and plan to take x if Tails (strategy 2); reject w and plan to take y if Tails (strategy 3). Their associated lotteries are $[w]$, $0.5[x] + 0.5[z]$, and $0.5[y] + 0.5[z]$ respectively.

By the stated preferences: $0.5[y] + 0.5[z] \succ [w] \succ 0.5[x] + 0.5[z]$, so strategy 3 delivers the best lottery.

(i) Can commit: The agent rejects w and commits to taking y if Tails. At the Tails node the commitment is honored even though $x \succ y$. Outcome: $0.5[y] + 0.5[z]$.

(ii) Cannot commit, foresees inconstancy: The agent knows that upon seeing Tails he will choose x over y (since $x \succ y$). So rejecting w actually delivers $0.5[x] + 0.5[z]$. Foreseeing this, the agent compares $[w]$ versus $0.5[x] + 0.5[z]$ and, since $[w] \succ 0.5[x] + 0.5[z]$, takes w. Outcome: $[w]$.

(iii) Cannot commit and cannot foresee: The agent naively rejects w, expecting to honor the plan of taking y at the Tails node. When Tails arrives he finds $x \succ y$ and switches to x. Outcome: $0.5[x] + 0.5[z]$, the worst of the three.

No behavioral assumption yields a coherent and consistent outcome without a commitment device. This is precisely why the substitution axioms – which rule out such preferences at the outset – are part of the definition of rationality.

(d) Solution. Regret theory proposes that agents experience disutility from not having chosen the ex-post better option, and that this additional term can rationalize the Allais pattern. From Myerson's perspective, regret-based explanations may describe psychological motivation, but they do not restore the behavioral consistency required by the substitution axioms. An agent who violates substitution faces the commitment problem of part (c) and will make choices he predictably regrets in a purely behavioral sense, regardless of any psychological story. Calling this regret-rationality does not resolve the underlying incoherence.

Problem 4 – Calculating Subjective Expected Utility

An entrepreneur chooses among three business strategies. The state space is $\Omega = \{t_H, t_M, t_L\}$, the prize space is $X = \{\$0, \$5, \$15, \$30\}$ (in thousands). Beliefs: $p(t_H) = 0.5$, $p(t_M) = 0.3$, $p(t_L) = 0.2$. Utility: $u(\$0) = 0$, $u(\$5) = 0.3$, $u(\$15) = 0.7$, $u(\$30) = 1$.

$$f = \begin{bmatrix} & \$0 & \$5 & \$15 & \$30 \\ t_H & 0 & 0 & 0.2 & 0.8 \\ t_M & 0.1 & 0.3 & 0.6 & 0 \\ t_L & 0.6 & 0.4 & 0 & 0 \end{bmatrix}, \quad g = \begin{bmatrix} & \$0 & \$5 & \$15 & \$30 \\ t_H & 0 & 0.1 & 0.5 & 0.4 \\ t_M & 0 & 0.4 & 0.6 & 0 \\ t_L & 0.3 & 0.7 & 0 & 0 \end{bmatrix}, \quad h = \begin{bmatrix} & \$0 & \$5 & \$15 & \$30 \\ t_H & 0 & 1 & 0 & 0 \\ t_M & 0 & 1 & 0 & 0 \\ t_L & 0 & 1 & 0 & 0 \end{bmatrix}$$

(a) Compute $E_p(u(f) \mid \Omega)$, $E_p(u(g) \mid \Omega)$, $E_p(u(h) \mid \Omega)$ and rank. **(b)** Condition on $S = \{t_H, t_M\}$, recompute beliefs via $p(t \mid S) = p(t)/p(S)$, recompute and re-rank. **(c)** Does the ranking change? Interpret economically. **(d)** Compute EMV for each act under prior beliefs and compare to the EU ranking.

(a) Solution. Using $E_p(u(f) \mid \Omega) = \sum_t p(t) \sum_x u(x)\, f(x \mid t)$.

First compute the conditional utility $V(f, t) = \sum_x u(x)\, f(x \mid t)$ in each state:

Strategy f:

- $V(f, t_H) = 0.7(0.2) + 1(0.8) = 0.14 + 0.80 = 0.94$
- $V(f, t_M) = 0(0.1) + 0.3(0.3) + 0.7(0.6) = 0 + 0.09 + 0.42 = 0.51$
- $V(f, t_L) = 0(0.6) + 0.3(0.4) = 0.12$

$$E_p(u(f) \mid \Omega) = 0.5(0.94) + 0.3(0.51) + 0.2(0.12) = 0.470 + 0.153 + 0.024 = \mathbf{0.647}$$

Strategy g:

- $V(g, t_H) = 0.3(0.1) + 0.7(0.5) + 1(0.4) = 0.03 + 0.35 + 0.40 = 0.78$
- $V(g, t_M) = 0.3(0.4) + 0.7(0.6) = 0.12 + 0.42 = 0.54$
- $V(g, t_L) = 0(0.3) + 0.3(0.7) = 0.21$

$$E_p(u(g) \mid \Omega) = 0.5(0.78) + 0.3(0.54) + 0.2(0.21) = 0.390 + 0.162 + 0.042 = \mathbf{0.594}$$

Strategy h: Since h gives \$5 with certainty in every state, $V(h, t) = u(\$5) = 0.3$ for all t.

$$E_p(u(h) \mid \Omega) = 0.5(0.3) + 0.3(0.3) + 0.2(0.3) = 0.3(1) = \mathbf{0.300}$$

Ranking: $f \succ g \succ h$. The entrepreneur should choose the aggressive launch f.

(b) Solution. Conditioning on $S = \{t_H, t_M\}$: $p(S) = 0.5 + 0.3 = 0.8$.

$$p(t_H \mid S) = \frac{0.5}{0.8} = 0.625, \qquad p(t_M \mid S) = \frac{0.3}{0.8} = 0.375.$$

$$E_p(u(f) \mid S) = 0.625(0.94) + 0.375(0.51) = 0.5875 + 0.19125 = \mathbf{0.779}$$

$$E_p(u(g) \mid S) = 0.625(0.78) + 0.375(0.54) = 0.4875 + 0.2025 = \mathbf{0.690}$$

$$E_p(u(h) \mid S) = 0.625(0.3) + 0.375(0.3) = 0.3(0.625 + 0.375) = \mathbf{0.300}$$

Conditional ranking: $f \succ g \succ h$. Unchanged.

(c) Solution. The ranking does not change. Conditioning on "demand is not low" reinforces the advantage of strategy f: it performs best precisely in the high-demand state t_H, which receives a larger weight (0.625 vs. 0.500) after ruling out t_L. If the ranking had changed, it would mean the new information shifts weight toward a state profile where a different strategy excels – i.e., the information would have decision-relevant content that alters optimal behavior. Here it does not.

(d) Solution. EMV $= \sum_t p(t) \sum_x x \cdot f(x \mid t)$ (prizes in thousands of dollars).

EMV of f:

- t_H: $15(0.2) + 30(0.8) = 3 + 24 = 27$
- t_M: $0(0.1) + 5(0.3) + 15(0.6) = 0 + 1.5 + 9 = 10.5$
- t_L: $0(0.6) + 5(0.4) = 2$

$$\text{EMV}(f) = 0.5(27) + 0.3(10.5) + 0.2(2) = 13.5 + 3.15 + 0.4 = \$\mathbf{17.05}\text{K}$$

EMV of g:

- t_H: $5(0.1) + 15(0.5) + 30(0.4) = 0.5 + 7.5 + 12 = 20$
- t_M: $5(0.4) + 15(0.6) = 2 + 9 = 11$
- t_L: $0(0.3) + 5(0.7) = 3.5$

$$\text{EMV}(g) = 0.5(20) + 0.3(11) + 0.2(3.5) = 10 + 3.3 + 0.7 = \$\mathbf{14.0}\text{K}$$

EMV of h: Gives \$5 in every state, so EMV(h) = \$**5**K.

EMV ranking: $f \succ g \succ h$ – the same as the EU ranking. The rankings agree in this instance, but this is not guaranteed in general. The utility function exhibits diminishing marginal utility ($u(\$15) - u(\$0) = 0.7$ but $u(\$30) - u(\$15) = 0.3$), reflecting risk aversion. A risk-averse agent penalizes variance relative to an EMV-maximizer, so if two acts had equal EMV but different risk profiles, their EU and EMV rankings could diverge. The agreement here is a numerical coincidence of the specific probability structure, not a general property of the model.

Problem 5 – Domination

Consider $X = \{\alpha, \beta, \gamma, \delta\}$ and $\Omega = \{\theta_1, \theta_2, \theta_3\}$ with payoff matrix:

Decision	θ_1	θ_2	θ_3
α	10	2	4
β	4	4	4
γ	2	6	5
δ	6	3	3

(a) For each decision find the set of beliefs $p = (p_1, p_2, p_3) \in \Delta(\Omega)$ under which it is optimal, and verify convexity. **(b)** Is δ strongly dominated? Find an explicit $\sigma \in \Delta(X)$ and verify. **(c)** Is β weakly or strongly dominated? Identify the dominating strategy. Is β ever a best response? **(d)** Verify Theorem 6 for δ directly: show no $p \in \Delta(\Omega)$ makes δ optimal.

(a) Solution. Let $p_1 + p_2 + p_3 = 1$, $p_i \geq 0$. Decision y is optimal iff $EU(y) \geq EU(x)$ for all x.

α **optimal.** The binding constraints are:

- vs. β: $10p_1 + 2p_2 + 4p_3 \geq 4p_1 + 4p_2 + 4p_3 \Rightarrow 6p_1 \geq 2p_2 \Rightarrow p_1 \geq \frac{1}{3}p_2$.
- vs. γ: $10p_1 + 2p_2 + 4p_3 \geq 2p_1 + 6p_2 + 5p_3 \Rightarrow 8p_1 \geq 4p_2 + p_3$.
- vs. δ: $10p_1 + 2p_2 + 4p_3 \geq 6p_1 + 3p_2 + 3p_3 \Rightarrow 4p_1 \geq p_2 - p_3$.

The set B_α is the intersection of these half-spaces with $\Delta(\Omega)$ – an intersection of convex sets, hence convex.

β **optimal.** The binding constraints are:

- vs. α: $4(p_1 + p_2 + p_3) \geq 10p_1 + 2p_2 + 4p_3 \Rightarrow 2p_2 \geq 6p_1 \Rightarrow p_2 \geq 3p_1$.
- vs. γ: $4(p_1 + p_2 + p_3) \geq 2p_1 + 6p_2 + 5p_3 \Rightarrow 2p_1 \geq 2p_2 + p_3$.

Combining: $2p_1 \geq 2p_2 + p_3 \geq 2(3p_1) + p_3 = 6p_1 + p_3$, giving $0 \geq 4p_1 + p_3$. Since $p_1, p_3 \geq 0$, this forces $p_1 = 0$ and $p_3 = 0$, hence $p_2 = 1$. But at $p = (0, 1, 0)$: $EU(\alpha) = 2$, $EU(\beta) = 4$, $EU(\gamma) = 6$, $EU(\delta) = 3$. Since $EU(\gamma) = 6 > 4 = EU(\beta)$, the constraint vs. γ is violated.

The constraint against δ is not needed, because the constraints against α and γ already imply impossibility.

Therefore **the set B_β is empty**: no belief makes β optimal. The empty set is convex.

γ **optimal.** The binding constraints are:

- vs. α: $2p_1 + 6p_2 + 5p_3 \geq 10p_1 + 2p_2 + 4p_3 \Rightarrow 4p_2 + p_3 \geq 8p_1$.
- vs. β: $2p_1 + 6p_2 + 5p_3 \geq 4(p_1 + p_2 + p_3) \Rightarrow 2p_2 + p_3 \geq 2p_1$.
- vs. δ: $2p_1 + 6p_2 + 5p_3 \geq 6p_1 + 3p_2 + 3p_3 \Rightarrow 3p_2 + 2p_3 \geq 4p_1$.

The set B_γ is again an intersection of half-spaces with $\Delta(\Omega)$ – convex. (For example, $p = (0, 1, 0)$ satisfies all three constraints, so B_γ is non-empty.)

δ **optimal.** The binding constraints are:

- vs. α: $6p_1 + 3p_2 + 3p_3 \geq 10p_1 + 2p_2 + 4p_3 \Rightarrow p_2 \geq 4p_1 + p_3$.
- vs. γ: $6p_1 + 3p_2 + 3p_3 \geq 2p_1 + 6p_2 + 5p_3 \Rightarrow 4p_1 \geq 3p_2 + 2p_3$.

From the first: $p_2 \geq 4p_1$. Substituting into the second: $4p_1 \geq 3p_2 \geq 12p_1$, giving $0 \geq 8p_1$, so $p_1 = 0$. Then $p_2 \geq p_3$ from the first constraint, and the second gives $0 \geq 3p_2 + 2p_3$, forcing $p_2 = p_3 = 0$ – contradicting $p_2 + p_3 = 1$. Therefore **the set B_δ is also empty**. The empty set is convex.

Summary: B_α and B_γ are non-empty convex sets; $B_\beta = B_\delta = \emptyset$.

(b) Solution. We seek $\sigma = \lambda[\alpha] + (1-\lambda)[\gamma]$ strongly dominating δ. The conditions are:

- θ_1: $10\lambda + 2(1-\lambda) > 6 \Rightarrow 8\lambda > 4 \Rightarrow \lambda > 0.5$.
- θ_2: $2\lambda + 6(1-\lambda) > 3 \Rightarrow 6 - 4\lambda > 3 \Rightarrow \lambda < 0.75$.
- θ_3: $4\lambda + 5(1-\lambda) > 3 \Rightarrow 5 - \lambda > 3 \Rightarrow \lambda < 2$. (always)

Any $\lambda \in (0.5, 0.75)$ works. Take $\lambda = 0.6$, so $\sigma = 0.6[\alpha] + 0.4[\gamma]$:

- θ_1: $0.6(10) + 0.4(2) = 6.8 > 6$. Yes.
- θ_2: $0.6(2) + 0.4(6) = 3.6 > 3$. Yes.
- θ_3: $0.6(4) + 0.4(5) = 4.4 > 3$. Yes.

δ is strongly dominated by $\sigma = 0.6[\alpha] + 0.4[\gamma]$.

(c) Solution. From part (a), $B_\beta = \emptyset$, so β is never a best response under any beliefs. To determine the type of domination, we look for $\sigma = \lambda[\alpha] + (1-\lambda)[\gamma]$ that dominates β, which pays exactly 4 in every state. The conditions for σ to beat β in each state are:

- θ_1: $8\lambda + 2 \geq 4 \Leftrightarrow \lambda \geq 0.25$.
- θ_2: $6 - 4\lambda \geq 4 \Leftrightarrow \lambda \leq 0.5$.
- θ_3: $5 - \lambda \geq 4 \Leftrightarrow \lambda \leq 1$. (always)

For $\lambda \in (0.25, 0.5)$, all three inequalities hold strictly: take $\lambda = 0.4$, so $\sigma = 0.4[\alpha] + 0.6[\gamma]$:

- θ_1: $0.4(10) + 0.6(2) = 5.2 > 4$. Yes.
- θ_2: $0.4(2) + 0.6(6) = 4.4 > 4$. Yes.
- θ_3: $0.4(4) + 0.6(5) = 4.6 > 4$. Yes.

Therefore β is **strongly dominated** (not merely weakly) by $\sigma = 0.4[\alpha] + 0.6[\gamma]$. It is never a best response under any $p \in \Delta(\Omega)$.

(d) Solution. For δ to be optimal, we need in particular $EU(\delta) \geq EU(\alpha)$ and $EU(\delta) \geq EU(\gamma)$.

vs. α:

$$6p_1 + 3p_2 + 3p_3 \geq 10p_1 + 2p_2 + 4p_3 \implies p_2 \geq 4p_1 + p_3. \tag{i}$$

vs. γ:

$$6p_1 + 3p_2 + 3p_3 \geq 2p_1 + 6p_2 + 5p_3 \implies 4p_1 \geq 3p_2 + 2p_3. \tag{ii}$$

Substitute (i) into (ii):

$$4p_1 \geq 3(4p_1 + p_3) + 2p_3 = 12p_1 + 5p_3 \implies 0 \geq 8p_1 + 5p_3.$$

Since $p_1, p_3 \geq 0$, we must have $p_1 = 0$ and $p_3 = 0$, hence $p_2 = 1$.

Now check whether δ is optimal at $p = (0, 1, 0)$: $EU(\delta) = 3$ and $EU(\beta) = 4$, so $EU(\beta) > EU(\delta)$. The constraint vs. β is violated.

Therefore no $p \in \Delta(\Omega)$ can make δ optimal, confirming the forward direction of Theorem 6 directly, without LP duality. ■

Problem 6 – Bayesian Conditional-Probability Systems and Zero-Probability Events

Let the state space be

$$\Omega = \{t_1, t_2, t_3\}.$$

For each $k = 1, 2, ...,$ define a full-support prior $\hat{p}^k \in \Delta^0(\Omega)$ by

$$\hat{p}^k(t_1) = 1 - \frac{1}{k} - \frac{1}{k^2}, \qquad \hat{p}^k(t_2) = \frac{1}{k}, \qquad \hat{p}^k(t_3) = \frac{1}{k^2}.$$

Let p be the limiting conditional-probability system generated by the sequence $\{\hat{p}^k\}_{k=1}^{\infty}$.

(a) Compute the limiting unconditional probabilities

$$p(t_1 \mid \Omega), \qquad p(t_2 \mid \Omega), \qquad p(t_3 \mid \Omega).$$

(b) Compute the limiting conditional probabilities after the event

$$S = \{t_2, t_3\}.$$

That is, compute

$$p(t_2 \mid S), \qquad p(t_3 \mid S).$$

(c) Compute the limiting conditional probabilities after the event

$$T = \{t_1, t_3\}.$$

That is, compute

$$p(t_1 \mid T), \qquad p(t_3 \mid T).$$

(d) Verify Bayes's formula for the nested events

$$R = \{t_3\} \subseteq S = \{t_2, t_3\} \subseteq \Omega.$$

That is, verify that

$$p(R \mid \Omega) = p(R \mid S)p(S \mid \Omega).$$

(e) Explain why this example shows that conditional beliefs after zero-probability events are not arbitrary. Even though $p(S \mid \Omega) = 0$, the conditional belief $p(t_2 \mid S)$ is still well-defined. What determines it?

(f) Think ahead to game theory. In a game, one player's unexpected action may reveal information or force other players to revise their beliefs. Why might it be useful to have a theory that assigns beliefs even after events that were initially considered extremely unlikely or impossible? Explain intuitively.

Solution.

For each $k = 1, 2, \ldots$, the prior is

$$\hat{p}^k(t_1) = 1 - \frac{1}{k} - \frac{1}{k^2}, \qquad \hat{p}^k(t_2) = \frac{1}{k}, \qquad \hat{p}^k(t_3) = \frac{1}{k^2}.$$

For large enough k, all three probabilities are strictly positive and sum to one:

$$\left(1 - \frac{1}{k} - \frac{1}{k^2}\right) + \frac{1}{k} + \frac{1}{k^2} = 1.$$

Thus $\hat{p}^k \in \Delta^0(\Omega)$ for large enough k. The limiting conditional-probability system p is obtained by conditioning $\hat{p}^k$ in the usual Bayesian way and then taking limits as $k \to \infty$.

$$p(t \mid S) = \lim_{k\to\infty} \frac{\hat{p}^k(t)}{\sum_{r\in S} \hat{p}^k(r)} \qquad \text{for } t \in S.$$

This is exactly the idea behind a Bayesian conditional-probability system: conditional beliefs can be understood as limits of ordinary Bayesian conditioning from full-support priors.

(a) Limiting unconditional probabilities.

Since $\Omega = \{t_1, t_2, t_3\}$, conditioning on Ω gives the unconditional prior itself:

$$p(t_i \mid \Omega) = \lim_{k\to\infty} \hat{p}^k(t_i).$$

For t_1:

$$p(t_1 \mid \Omega) = \lim_{k\to\infty} \left(1 - \frac{1}{k} - \frac{1}{k^2}\right) = 1.$$

For t_2:

$$p(t_2 \mid \Omega) = \lim_{k\to\infty} \frac{1}{k} = 0.$$

For t_3:

$$p(t_3 \mid \Omega) = \lim_{k\to\infty} \frac{1}{k^2} = 0.$$

Therefore,

$$\boxed{p(t_1 \mid \Omega) = 1, \qquad p(t_2 \mid \Omega) = 0, \qquad p(t_3 \mid \Omega) = 0.}$$

So in the limiting system, state t_1 receives probability one, while t_2 and t_3 are zero-probability states.

(b) Conditional probabilities after $S = \{t_2, t_3\}$.

Now condition on the event

$$S = \{t_2, t_3\}.$$

For each finite k, the conditional probability of t_2 given S is

$$\hat{p}^k(t_2 \mid S) = \frac{\hat{p}^k(t_2)}{\hat{p}^k(t_2) + \hat{p}^k(t_3)}.$$

Substitute the probabilities:

$$\hat{p}^k(t_2 \mid S) = \frac{\frac{1}{k}}{\frac{1}{k} + \frac{1}{k^2}}.$$

Simplify by multiplying numerator and denominator by k^2:

$$\hat{p}^k(t_2 \mid S) = \frac{k}{k+1}.$$

Therefore,

$$p(t_2 \mid S) = \lim_{k\to\infty} \frac{k}{k+1} = 1.$$

Similarly,

$$\hat{p}^k(t_3 \mid S) = \frac{\hat{p}^k(t_3)}{\hat{p}^k(t_2) + \hat{p}^k(t_3)} = \frac{\frac{1}{k^2}}{\frac{1}{k} + \frac{1}{k^2}}.$$

Multiplying numerator and denominator by k^2 gives

$$\hat{p}^k(t_3 \mid S) = \frac{1}{k+1}.$$

Thus,

$$p(t_3 \mid S) = \lim_{k\to\infty} \frac{1}{k+1} = 0.$$

Therefore,

$$\boxed{p(t_2 \mid S) = 1, \qquad p(t_3 \mid S) = 0.}$$

This is the key point of the example. Although both t_2 and t_3 have zero limiting probability under Ω, they do not vanish at the same rate. State t_2 has probability order $1/k$, while state t_3 has probability order $1/k^2$. Since $1/k$ is much larger than $1/k^2$ for large k, conditional on the rare event $\{t_2, t_3\}$, the limiting belief assigns probability one to t_2.

(c) Conditional probabilities after $T = \{t_1, t_3\}$.

Now condition on

$$T = \{t_1, t_3\}.$$

For each k,

$$\hat{p}^k(t_1 \mid T) = \frac{\hat{p}^k(t_1)}{\hat{p}^k(t_1) + \hat{p}^k(t_3)}.$$

Substitute:

$$\hat{p}^k(t_1 \mid T) = \frac{1 - \frac{1}{k} - \frac{1}{k^2}}{1 - \frac{1}{k} - \frac{1}{k^2} + \frac{1}{k^2}}.$$

The denominator simplifies to

$$1 - \frac{1}{k}.$$

So

$$\hat{p}^k(t_1 \mid T) = \frac{1 - \frac{1}{k} - \frac{1}{k^2}}{1 - \frac{1}{k}}.$$

Taking limits,

$$p(t_1 \mid T) = \lim_{k \to \infty} \frac{1 - \frac{1}{k} - \frac{1}{k^2}}{1 - \frac{1}{k}} = 1.$$

For t_3,

$$\hat{p}^k(t_3 \mid T) = \frac{\hat{p}^k(t_3)}{\hat{p}^k(t_1) + \hat{p}^k(t_3)} = \frac{\frac{1}{k^2}}{1 - \frac{1}{k}}.$$

Thus,

$$p(t_3 \mid T) = \lim_{k \to \infty} \frac{\frac{1}{k^2}}{1 - \frac{1}{k}} = 0.$$

Therefore,

$$\boxed{p(t_1 \mid T) = 1, \qquad p(t_3 \mid T) = 0.}$$

This result is intuitive because t_1 remains overwhelmingly likely relative to t_3.

(d) Verification of Bayes's formula.

We need to verify Bayes's formula for

$$R = \{t_3\} \subseteq S = \{t_2, t_3\} \subseteq \Omega.$$

The formula is

$$p(R \mid \Omega) = p(R \mid S)p(S \mid \Omega).$$

First compute the left-hand side. Since $R = \{t_3\}$,

$$p(R \mid \Omega) = p(t_3 \mid \Omega) = 0.$$

Now compute the right-hand side. From part (b),

$$p(R \mid S) = p(t_3 \mid S) = 0.$$

Also,

$$p(S \mid \Omega) = p(\{t_2, t_3\} \mid \Omega) = p(t_2 \mid \Omega) + p(t_3 \mid \Omega) = 0 + 0 = 0.$$

Therefore,

$$p(R \mid S)p(S \mid \Omega) = 0 \cdot 0 = 0.$$

Hence,

$$p(R \mid \Omega) = p(R \mid S)p(S \mid \Omega),$$

because both sides equal zero:

$$\boxed{0 = 0.}$$

Bayes's formula is satisfied.

(e) Why zero-probability conditional beliefs are not arbitrary.

The event

$$S = \{t_2, t_3\}$$

has limiting probability zero:

$$p(S \mid \Omega) = p(t_2 \mid \Omega) + p(t_3 \mid \Omega) = 0 + 0 = 0.$$

However, the conditional belief after S is still well-defined:

$$p(t_2 \mid S) = 1, \qquad p(t_3 \mid S) = 0.$$

This may seem surprising at first. In ordinary probability theory, conditioning on a zero-probability event is usually undefined. But a Bayesian conditional-probability system adds more structure. It asks: if the event S is reached as the limit of small but positive probability events, what are the relative probabilities of states inside S?

Here, before taking the limit,

$$\hat{p}^k(t_2) = \frac{1}{k}, \qquad \hat{p}^k(t_3) = \frac{1}{k^2}.$$

Both probabilities go to zero, but t_3 goes to zero faster. The ratio is

$$\frac{\hat{p}^k(t_3)}{\hat{p}^k(t_2)} = \frac{\frac{1}{k^2}}{\frac{1}{k}} = \frac{1}{k} \to 0.$$

So, conditional on the rare event $\{t_2, t_3\}$, state t_2 is infinitely more likely than state t_3 in the limit. This is why

$$p(t_2 \mid S) = 1 \quad \text{and} \quad p(t_3 \mid S) = 0.$$

Thus, zero-probability conditional beliefs are not arbitrary. They are determined by the relative rates at which the probabilities of the states go to zero.

(f) Why this idea matters in game theory.

In game theory, players often have to think about what they would believe after observing another player's action. Some actions may be surprising: according to the theory, they were not supposed to happen, or they were considered extremely unlikely.

Even then, the game does not simply stop. The other players must still ask: if this surprising action occurred, what should I now believe? Which state, type, or intention is more likely?

This example shows why conditional beliefs after unlikely events are useful. Even when an event has limiting probability zero, the theory can still give a disciplined answer about what is more likely inside that event.

Here,

$$p(S \mid \Omega) = 0,$$

but conditional on $S = \{t_2, t_3\}$, the theory still tells us

$$p(t_2 \mid S) = 1 \quad \text{and} \quad p(t_3 \mid S) = 0.$$

So the main lesson is: in game theory, we often need beliefs not only along the expected path of play, but also after surprising events. Conditional-probability systems provide a way to make those beliefs consistent rather than arbitrary.