

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

Problem Set 1 - Basic Models of Game Theory

Instructions: You are encouraged to solve the problems before the recitation. Additionally, you are encouraged to work in groups. It is **not mandatory** to submit solutions unless stated otherwise. However, if you would like to share your solution, I would be happy to review it.

Problem 1 We consider a situation involving a **Firm** and a **Worker**.

- The **Worker** can be of two types:
 - **High ability** (prefers to Work if hired)
 - **Low ability** (prefers to Shirk if hired)
- The **Firm** wants to hire the Worker **only if he will Work**.
- The Worker **knows his own type**, but the Firm does **not**.
- The Firm knows that the Worker knows his own type.

Let:

- Probability that the Worker is **High ability**: p
- Probability that the Worker is **Low ability**: $1 - p$

If hired, the Worker chooses whether to **Work** or **Shirk**.

The resulting **payoffs (Firm, Worker)** are:

High Ability Worker

- Hire & Work: $(1, 2)$
- Hire & Shirk: $(0, 1)$
- Not Hire: $(0, 0)$

Low Ability Worker

- Hire & Work: $(1, 1)$
- Hire & Shirk: $(-1, 2)$
- Not Hire: $(0, 0)$

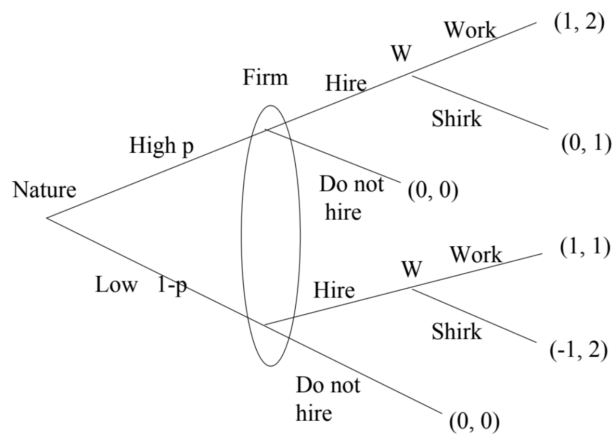
(a) Extensive Form

- Write this situation as an **extensive-form game** with imperfect information.
- Carefully label:
 - The sequence of moves (Nature, Worker, Firm)
 - Information sets
 - Actions and terminal payoffs

Solution:

- **Nature** moves first and chooses the Worker's type:
 - High type with probability p
 - Low type with probability $1 - p$
- The Worker observes their type.
- The **Firm**, without knowing the type, decides:
 - Hire (H) or Not Hire (N)
- If Hired, the Worker chooses:
 - Work (W) or Shirk (S)

There are two information sets for the Worker, depending on the type, and a single information set for the Firm.



(b) Normal Form

- Rewrite the problem in **normal (strategic) form**:
 - Enumerate the **strategy sets** for each player.
 - Construct the **payoff matrix** by considering all strategy combinations.

Solution:

Let:

- Firm: H = Hire, NH = Not Hire
- Worker strategies:
 - WW = (High \rightarrow Work, Low \rightarrow Work)
 - WS = (High \rightarrow Work, Low \rightarrow Shirk)

– SW = (High → Shirk, Low → Work)

– SS = (High → Shirk, Low → Shirk)

Firm / Worker	WW	WS	SW	SS
H	(1, 1 + p)	(2p - 1, 2)	(1 - p, 1)	(-1 + p, 2 - p)
NH	(0, 0)	(0, 0)	(0, 0)	(0, 0)

(c) Multi-Agent Representation

- Represent the situation using a **multi-agent normal form**:
 - Treat each **type of Worker** as a separate agent.
 - List all possible strategies and assign payoffs.

Solution:

	Firm = Hire		Firm = Not hire	
	$W_L = W$	$W_L = S$	$W_L = W$	$W_L = S$
$W_H = W$	(1, 1+p, 1+p)	(2p-1, 2, 2)	(0, 0, 0)	(0, 0, 0)
$W_H = S$	(1-p, 1, 1)	(-1+p, 2-p, 2-p)	(0, 0, 0)	(0, 0, 0)

(d) Bayesian Game Representation

- Model the situation as a **Bayesian game**:
 - Define:
 - * The set of players
 - * The set of types
 - * Strategy sets
 - * Beliefs
 - * Payoff functions
 - Describe how **incomplete information** is handled through beliefs and types.

Solution:

One can write the game in the example above as a Bayesian game by setting

- $N = \{F, W\}$
- $T_F = \{t_F\}$, $T_W = \{\text{high, low}\}$
- $p(t_F, \text{high}) = p$, $p(t_F, \text{low}) = 1 - p$
- $A_F = \{\text{Hire, Not hire}\}$, $A_W = \{\text{Work, Shirk}\}$

and the utility functions u_F and u_W are defined by the following tables, where the first entry is the payoff of the firm and the second the payoff of the worker. The table on the left corresponds to type $t = (t_F, \text{high})$ and the one on the right to $t = (t_F, \text{low})$:

$t_W = \text{High}$	Work	Shirk	$t_W = \text{low}$	Work	Shirk
Hire	(1, 2)	(0, 1)	Hire	(1, 1)	(-1, 2)
Not hire	(0, 0)	(0, 0)	Not Hire	(0, 0)	(0, 0)

Problem 2 Suppose that strategy s_1 strictly dominates strategy s'_1 , and that strategy s_2 strictly dominates strategy s'_2 . Let σ_1 be a mixture that attaches positive probability to (only) s_1 and s_2 , and σ'_1 a mixture that attaches positive probability to (only) s'_1 and s'_2 . Either prove that σ_1 strictly dominates σ'_1 , or provide a counterexample.

Solution:

It suffices to show a counterexample to prove that it may be the case that σ_1 does not strictly dominate σ'_1 . Consider the following game (showing only player 1's payoffs),

	L	R
s_1	2	5
s'_1	1	4
s_2	1	2
s'_2	0	1

It is easy to see that,

$$u_1(s_1, s) > u_1(s'_1, s) \quad \text{for all } s \in S_2$$

$$u_1(s_2, s) > u_1(s'_2, s) \quad \text{for all } s \in S_2$$

but, for $\sigma_1 = (\frac{1}{5}, 0, \frac{4}{5}, 0)$, $\sigma'_1 = (0, \frac{7}{8}, 0, \frac{1}{8})$ we have

$$u_1(\sigma_1, R) = \frac{13}{5} < \frac{29}{8} = u_1(\sigma'_1, R).$$

Problem 3 Consider a game in which player 1 names an amount of money x from the interval $[0, 100]$ to be given to player 2. Player 2 simultaneously names

the minimum amount of money y that 2 is willing to accept. If $y > x$, neither player receives anything. If $y \leq x$, player 2 receives x and player 1 receives $100 - x$. Which strategies survive the iterated elimination of weakly dominated strategies? Now consider the same game, except that if $y < x$, 2 receives y and 1 receives $100 - y$. Now which strategies survive iterated weak dominance? Which payoff rule seems to favor player 2? Under which does 2 fare better? (Assume that you can equate monetary payoffs with utilities.)

Solution:

Strategies spaces are $S_1^0 = [0, 100]$, and $S_2^0 = [0, \infty)$, and payoffs are given by

$$\begin{aligned} u_1(x, y) &= 1(y \leq x)(100 - x), & x \in [0, 100] \\ u_2(x, y) &= 1(y \leq x)x, & x \in [0, 100] \end{aligned}$$

First, note that for player 2 any $y > 100$ is weakly dominated by 100, but we can go further and note that any $y > 0$ is weakly dominated by 0, i.e., $u_2(x, y) \leq u_2(x, 0)$, $\forall x, \forall y > 0$, with strict inequality for some x . To see this, take $\tilde{y} > 0$, if $x \geq \tilde{y}$ (or $x = 0$) then $u_2(x, \tilde{y}) = u_2(x, 0) = x$, but for $0 < x < \tilde{y}$, $u_2(x, \tilde{y}) = 0 < x = u_2(x, 0)$. Hence $S_2^1 = \{0\}$.

Now, having eliminated all $y > 0$ for player 2, all $x > 0$ are weakly (in fact, strictly) dominated by 0 for player 1 (since $u_1(x, 0) = 100 - x$, $x \in [0, 100]$, is decreasing in x). Hence $S_1^2 = \{0\}$. Therefore, $S_i^* \equiv \bigcap_{n=0}^{\infty} S_i^n = \{0\}$, $\forall i$.

In the second formulation of the game we have $S_1^0 = [0, 100]$, and $S_2^0 = [0, \infty)$, and payoffs given by

$$\begin{aligned} u_1(x, y) &= 1(y \leq x)(100 - y), & x \in [0, 100] \\ u_2(x, y) &= 1(y \leq x)y, & x \in [0, 100] \end{aligned}$$

Note that, for player 1, any $x < 100$ is weakly dominated by 100, since for any $x < 100$ and all $y \geq 0$, $u_1(x, y) \leq u_1(100, y)$, but for y such that $x < y < 100$, $u_1(x, y) = 0 < 100 - y = u_1(100, y)$. Hence $S_1^1 = \{100\}$, $S_2^1 = S_2^0$.

Now, having eliminated all $x < 100$ for player 1, all $y \neq 100$ are weakly (in fact, strictly) dominated by 100 for player 2 (since $u_2(100, y) = y$, $y \in [0, 100]$, is increasing in y , and $u_2(100, y) = 0$ for $y > 100$). Hence $S_2^2 = \{100\}$. Therefore, $S_i^* \equiv \bigcap_{n=0}^{\infty} S_i^n = \{100\}$, $\forall i$.

In the first game the resulting payoffs are $u_1(0, 0) = 100$ and $u_2(0, 0) = 0$, while in the second $u_1(100, 100) = 0$ and $u_2(100, 100) = 100$. Therefore, in the second rule player 2 fares better.

Problem 4 Consider a three-player extensive-form game. Nature first draws a state $\omega \in \{A, B\}$, each with probability $\frac{1}{2}$. Player 1 observes ω and chooses T or B . Player 2 observes only Player 1's action (not ω) and chooses L or R . Player 3 observes neither ω nor Player 1's action, but does observe Player 2's action, and chooses U or D . Payoffs (u_1, u_2, u_3) at terminal nodes are:

$\omega = A$	L	R	$\omega = B$	L	R
T, U	(3, 2, 1)	(0, 1, 2)	T, U	(1, 0, 2)	(3, 1, 0)
T, D	(1, 3, 0)	(2, 0, 3)	T, D	(2, 3, 1)	(0, 2, 3)
B, U	(0, 1, 3)	(3, 2, 0)	B, U	(3, 2, 0)	(1, 0, 1)
B, D	(2, 0, 2)	(1, 3, 1)	B, D	(0, 1, 3)	(2, 3, 2)

(a) Draw the extensive-form game carefully, labeling all nodes, information sets, and terminal payoffs. How many information sets does each player have? Verify that perfect recall holds for all three players.

(b) Write down the strategy sets C_1, C_2, C_3 . How many pure strategy profiles are there in total?

(c) Construct the multiagent representation. How many agents are there, and what is each agent's strategy set? Explain why separating a player into temporary agents changes how deviations are checked. What does the multiagent representation help us see, and what does it not determine by itself?

(d) Write the game in Bayesian form Γ^b . Define the type sets, strategy functions, and utility functions formally. Player 2's belief over ω is derived from observing Player 1's action – explain why this belief is in general not equal to the prior $\frac{1}{2}$, and what additional structure would be needed to pin it down.

(e) Suppose we restrict attention to strategies in which Player 1 plays the same action regardless of ω (a **pooling strategy**). Show that under any pooling strategy by Player 1, Player 2's information set is uninformative about ω . Analyze Player 3's situation separately under pooling on T and pooling on B : in each case, identify Player 3's optimal contingent strategy and determine whether a strictly dominant action exists at both information sets.

Solution

Part (a): Extensive Form

Tree structure.

- **Nature** moves at the root, choosing $\omega \in \{A, B\}$ with equal probability $\frac{1}{2}$.
- **Player 1** observes ω and moves at two nodes, one for each realization of ω . Since Player 1 knows ω , each node is a **singleton information set**.
- **Player 2** observes Player 1's action but not ω . She therefore has two information sets:

$$S_2 = \{s_T, s_B\}$$

where s_T is reached when Player 1 plays T (containing two nodes, one from $\omega = A$ and one from $\omega = B$), and s_B is reached when Player 1 plays B (similarly two nodes).

- **Player 3** observes Player 2's action but neither ω nor Player 1's action. She therefore has two information sets:

$$S_3 = \{s_L, s_R\}$$

where s_L is reached when Player 2 plays L (containing four nodes, one for each combination $(\omega, a_1) \in \{A, B\} \times \{T, B\}$ leading to L), and s_R is reached when Player 2 plays R (four nodes similarly).

Information set count:

- Player 1: **2 information sets** (singleton, since Player 1 observes ω)
- Player 2: **2 information sets** (s_T and s_B , each containing 2 nodes)
- Player 3: **2 information sets** (s_L and s_R , each containing 4 nodes)

Verification of perfect recall.

Perfect recall requires that each player remembers all previous moves they made and all previous information states they passed through.

In this game, no strategic player moves more than once along any realized path of play. Player 1 moves once (at her ω -specific node), Player 2 moves once (at s_T or s_B), and Player 3 moves once (at s_L or s_R). Since no player has a second information state to reach after their first move, there is no earlier own action or information state that could be forgotten. Perfect recall therefore holds trivially for all three players.

Part (b): Strategy Sets

A strategy must specify an action at **every** information state, including those that may not be reached in actual play. This is a requirement of the definition: a strategy is a complete contingent plan.

Player 1 has 2 information states ($\omega = A$ and $\omega = B$), each with actions $\{T, B\}$:

$$C_1 = \{TT, TB, BT, BB\}$$

where the first letter is the action at $\omega = A$ and the second at $\omega = B$.

Player 2 has 2 information states (s_T and s_B), each with actions $\{L, R\}$:

$$C_2 = \{LL, LR, RL, RR\}$$

where the first letter is the action at s_T and the second at s_B .

Player 3 has 2 information states (s_L and s_R), each with actions $\{U, D\}$:

$$C_3 = \{UU, UD, DU, DD\}$$

where the first letter is the action at s_L and the second at s_R .

This is exactly the product structure from the definition of strategies:

$$C_i = \prod_{s \in S_i} D_s$$

Total pure strategy profiles:

$$|C_1| \times |C_2| \times |C_3| = 4 \times 4 \times 4 = 64$$

Part (c): Multiagent Representation

Agents. Each information state of each player becomes a separate temporary agent:

- Player 1 generates agents 1_A (moves at $\omega = A$) and 1_B (moves at $\omega = B$)
- Player 2 generates agents 2_T (moves at s_T) and 2_B (moves at s_B)
- Player 3 generates agents 3_L (moves at s_L) and 3_R (moves at s_R)

Total: **6 temporary agents**, each sharing the preferences of their original player. Their strategy sets are:

$$D_{1_A} = D_{1_B} = \{T, B\}, \quad D_{2_T} = D_{2_B} = \{L, R\}, \quad D_{3_L} = D_{3_R} = \{U, D\}$$

Why separating a player into temporary agents changes how deviations are checked.

In the original normal-form game, Player 1 chooses a complete contingent plan from $C_1 = \{TT, TB, BT, BB\}$. When we ask whether Player 1 has a profitable deviation, we must compare her entire plan against every other entire plan. A deviation may require changing the action at both information states simultaneously – for example, switching from TB to BT changes the action at both $\omega = A$ and $\omega = B$ at once.

In the multiagent representation, 1_A and 1_B are separate agents. Agent 1_A can only change the action at $\omega = A$, holding 1_B 's action fixed. Agent 1_B can only change the action at $\omega = B$, holding 1_A 's action fixed. The deviation check is therefore strictly more limited: we only ask whether a single agent, at a single information state, wishes to change her action. A coordinated change across two agents of the same original player is never considered, because no single agent initiates it.

This means a strategy profile that would be challenged by a joint deviation in the original normal form may go unchallenged in the multiagent representation, simply because neither individual agent finds the one-sided change profitable on its own. The multiagent representation therefore imposes a weaker discipline on strategy profiles than the original normal form.

What the multiagent representation helps us see, and what it does not determine.

The multiagent representation separates decisions by information state and makes it possible to ask whether behavior is optimal at each node independently. This is the foundation for reasoning information-set by information-set: rather than asking whether a player's overall plan is optimal, we ask whether the action at each individual information state is optimal given what the player knows there.

However, separating decisions by information state is only part of what this requires. We also need to specify what a player *believes* at each information state – in particular at information states that are reached with zero probability under the strategies being played. The multiagent representation does not resolve this belief problem. It separates the decision points, but beliefs must still be specified separately when Bayes' rule does not determine them.

Part (d): Bayesian Form

The Bayesian game is:

$$\Gamma^b = (N, (C_i)_{i \in N}, (T_i)_{i \in N}, (p_i)_{i \in N}, (u_i)_{i \in N})$$

Players: $N = \{1, 2, 3\}$.

Type sets. Player 1's private information is ω :

$$T_1 = \{A, B\}$$

Players 2 and 3 have no initial private type:

$$T_2 = \{t_2\}, \quad T_3 = \{t_3\}$$

Primitive action sets:

$$C_1 = \{T, B\}, \quad C_2 = \{L, R\}, \quad C_3 = \{U, D\}$$

Strategy functions. Because the game is dynamic, Players 2 and 3 also condition on what they observe during play. The relevant behavioral rules are therefore:

$$s_1 : \{A, B\} \rightarrow \{T, B\}$$

$$s_2 : \{T, B\} \rightarrow \{L, R\}$$

$$s_3 : \{L, R\} \rightarrow \{U, D\}$$

Here s_2 conditions on Player 1's observed action, and s_3 conditions on Player 2's observed action – not on private types, since Players 2 and 3 have none.

Utility functions: $u_i : C \times T_1 \rightarrow \mathbb{R}$, read directly from the payoff tables, where the relevant state is $\omega \in T_1$.

Why Player 2's posterior differs from the prior.

Player 2 observes $a_1 \in \{T, B\}$ before choosing. Her posterior belief about ω is derived via Bayes' rule from Player 1's strategy. Let $\sigma_1(a_1 | \omega)$ denote the probability that Player 1 plays a_1 given type ω . Then:

$$\mu_2(A | a_1) = \frac{\sigma_1(a_1 | A) \cdot \frac{1}{2}}{\sigma_1(a_1 | A) \cdot \frac{1}{2} + \sigma_1(a_1 | B) \cdot \frac{1}{2}} = \frac{\sigma_1(a_1 | A)}{\sigma_1(a_1 | A) + \sigma_1(a_1 | B)}$$

This equals the prior $\frac{1}{2}$ only when $\sigma_1(a_1 | A) = \sigma_1(a_1 | B)$, i.e., when Player 1's strategy is uninformative about ω – the pooling case. In general, whenever Player 1's strategy is even partially separating, the posterior differs from the prior.

When the observed action a_1 is not reached under the proposed behavior, so that

$$\sigma_1(a_1 | A) = \sigma_1(a_1 | B) = 0,$$

Bayes' rule does not apply and the posterior is not pinned down. In that case, one must specify the belief directly.

Part (e): Pooling Strategies and Player 3's Optimal Actions

Claim 1. Under any pooling strategy by Player 1, Player 2's information set is uninformative about ω .

Proof. Suppose $s_1(A) = s_1(B) = a^*$. Then $\sigma_1(a^* | A) = \sigma_1(a^* | B) = 1$. Applying Bayes' rule at the reached information set:

$$\mu_2(A | a^*) = \frac{1 \cdot \frac{1}{2}}{1 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2}} = \frac{1}{2}$$

The posterior equals the prior. Observing a^* provides no information about ω . \square

What Player 3's information set reveals under pooling.

Under pooling by Player 1, Player 2 cannot distinguish ω at her on-path information set. Her action a_2^* therefore carries no information about ω . Player 3 observes a_2^* , which itself was taken without knowledge of ω . Therefore Player 3's posterior also satisfies $p_3(\omega = A | a_2) = \frac{1}{2}$ – the entire history of play is uninformative about ω under pooling.

Analysis of Player 3's optimal contingent strategy.

We analyze both pooling cases separately, as the answer differs.

Case 1: Player 1 pools on T , i.e., $s_1(A) = s_1(B) = T$.

Player 3 holds posterior $p_3(\omega = A) = \frac{1}{2}$ at both information sets.

At information set s_L (after Player 2 plays L):

$$\mathbb{E}[u_3 \mid T, L, U] = \frac{1}{2}(1) + \frac{1}{2}(2) = 1.5$$

$$\mathbb{E}[u_3 \mid T, L, D] = \frac{1}{2}(0) + \frac{1}{2}(1) = 0.5$$

Player 3 strictly prefers U after L .

At information set s_R (after Player 2 plays R):

$$\mathbb{E}[u_3 \mid T, R, U] = \frac{1}{2}(2) + \frac{1}{2}(0) = 1$$

$$\mathbb{E}[u_3 \mid T, R, D] = \frac{1}{2}(3) + \frac{1}{2}(3) = 3$$

Player 3 strictly prefers D after R .

Conclusion for Case 1. Player 3's optimal contingent strategy is UD . Since the strictly preferred action differs across information sets – U after L and D after R – **no single action is strictly dominant** across both information sets under pooling on T .

Case 2: Player 1 pools on B , i.e., $s_1(A) = s_1(B) = B$.

Player 3 again holds posterior $p_3(\omega = A) = \frac{1}{2}$ at both information sets.

At information set s_L (after Player 2 plays L):

$$\mathbb{E}[u_3 \mid B, L, U] = \frac{1}{2}(3) + \frac{1}{2}(0) = 1.5$$

$$\mathbb{E}[u_3 \mid B, L, D] = \frac{1}{2}(2) + \frac{1}{2}(3) = 2.5$$

Player 3 strictly prefers D after L .

At information set s_R (after Player 2 plays R):

$$\mathbb{E}[u_3 \mid B, R, U] = \frac{1}{2}(0) + \frac{1}{2}(1) = 0.5$$

$$\mathbb{E}[u_3 \mid B, R, D] = \frac{1}{2}(1) + \frac{1}{2}(2) = 1.5$$

Player 3 strictly prefers D after R .

Conclusion for Case 2. Player 3's optimal contingent strategy is DD . Since D is strictly preferred at both s_L and s_R , the action D is **strictly dominant** for Player 3 under pooling on B .

Overall conclusion. Whether Player 3 has a strictly dominant action depends on which action Player 1 pools on:

- Under **pooling on T** : optimal contingent strategy is UD ; no single action is strictly dominant across both information sets.
- Under **pooling on B** : optimal contingent strategy is DD ; action D is strictly dominant at both information sets.

The information structure is identical in both cases – Player 3 always holds the prior $\frac{1}{2}$ – but the payoff consequences of the pooled action differ, which changes the dominance structure Player 3 faces.

Problem 5 A government (Player G) must decide whether to **Audit** (A) or **Not Audit** (N) a firm (Player F). The firm has private information about its type: it is either **Compliant** (C) or **Non-Compliant** (NC), with probabilities p and $1 - p$ respectively.

If the government audits, it incurs cost $c > 0$. If it audits a non-compliant firm, it collects fine $f > c$. If it audits a compliant firm, it collects nothing but pays c . A non-compliant firm that is not audited gains $g > 0$.

	A	N
C	$(-c, 0)$	$(0, 0)$
NC	$(f - c, -f)$	$(0, g)$

where rows are firm types, columns are government actions, and payoffs are (Government, Firm). The firm first chooses **Disclose** (D) or **Conceal** (K), and the government observes this disclosure decision before deciding whether to audit.

- Write this as a Bayesian game Γ^b . Define all components formally.
- Draw the extensive form of the extended game. How many information sets does the government have? Write down the government's strategy set.
- With $p = \frac{2}{3}$, $c = 1$, $f = 4$, $g = 3$: for what values of μ (the government's posterior belief that the firm is non-compliant, given its observed message) does the government strictly prefer to audit?
- Under what conditions does a non-compliant firm prefer to mimic the compliant type by playing D ? If both types pool on D , what is the posterior $\mu(D)$? Is auditing rational at this posterior? Does the non-compliant firm strictly prefer to remain at D , or only weakly prefer it? Explain how the answer depends on the government's belief after the message K .
- Suppose the compliant type plays D with probability 1, and the non-compliant type plays D with probability $\lambda \in [0, 1]$ and K with probability $1 - \lambda$. The government audits with probability α after D and β after K . Find values of α , β , and λ such that the government's audit decisions are consistent with its posterior beliefs and the non-compliant firm is indifferent between D and K . Under what parameter restriction does an interior value $\lambda \in (0, 1)$ exist?

Solution

Part (a): Bayesian Game

$$\Gamma^b = (N, (C_i)_{i \in N}, (T_i)_{i \in N}, (p_i)_{i \in N}, (u_i)_{i \in N})$$

Players: $N = \{G, F\}$.

Type sets:

$$T_F = \{C, NC\}, \quad T_G = \{t_G\}$$

The firm has private type $t_F \in T_F$; the government has no private information.

Action sets:

$$C_F = \{D, K\}, \quad C_G = \{A, N\}$$

Strategies. The firm's strategy is a function mapping each type to a disclosure decision:

$$s_F : T_F \rightarrow C_F$$

So the firm's strategy set is:

$$S_F = \{DD, DK, KD, KK\}$$

where the first letter is the action of type C and the second is the action of type NC .

The government observes the firm's message $m \in \{D, K\}$ and conditions its audit decision on it:

$$s_G : \{D, K\} \rightarrow \{A, N\}$$

So the government's strategy set is:

$$S_G = \{AA, AN, NA, NN\}$$

where the first letter is the action after D and the second after K .

Beliefs. The government's prior over the firm's type is:

$$\Pr(NC) = 1 - p, \quad \Pr(C) = p$$

After observing message $m \in \{D, K\}$, the government holds a posterior belief:

$$\mu_G(t_F | m) \in \Delta(\{C, NC\})$$

derived via Bayes' rule from the firm's strategy. We write $\mu_G(NC | m)$ as $\mu(m)$ for brevity.

Utility functions. For the compliant firm:

$$u_F(m, A, C) = 0, \quad u_F(m, N, C) = 0 \quad \forall m$$

For the non-compliant firm:

$$u_F(m, A, NC) = -f, \quad u_F(m, N, NC) = g \quad \forall m$$

For the government:

$$u_G(m, A, C) = -c, \quad u_G(m, A, NC) = f - c, \quad u_G(m, N, \cdot) = 0 \quad \forall m$$

Part (b): Extensive Form

The tree proceeds as follows:

1. **Nature** draws $t_F \in \{C, NC\}$ with probabilities $(p, 1 - p)$
2. **Firm** observes t_F and chooses $m \in \{D, K\}$
3. **Government** observes m (but not t_F) and chooses $a_G \in \{A, N\}$
4. Terminal nodes assign payoffs (u_G, u_F)

Government's information sets. The government has **2 information sets**:

- One reached after observing D , containing two nodes: (C, D) and (NC, D)
- One reached after observing K , containing two nodes: (C, K) and (NC, K)

Government's strategy set:

$$S_G = \{AA, AN, NA, NN\}$$

Part (c): Government's Audit Decision

Let $\mu \in [0, 1]$ be the government's posterior belief that the firm is non-compliant, given the observed message. The government's expected payoffs are:

$$\mathbb{E}[u_G \mid A, \mu] = \mu(f - c) + (1 - \mu)(-c) = \mu f - c$$

$$\mathbb{E}[u_G \mid N, \mu] = 0$$

The government strictly prefers to audit iff:

$$\mu f - c > 0 \iff \mu > \frac{c}{f}$$

With $c = 1$ and $f = 4$:

$$\mu > \frac{1}{4}$$

Interpretation. The government audits when the expected fine revenue μf exceeds the audit cost c . The threshold $\frac{c}{f}$ depends only on the cost-to-fine ratio, not on p or g .

Part (d): Pooling on D – Payoff Analysis

When does the non-compliant firm prefer D ?

Let $\alpha = \Pr(A \mid D)$ and $\beta = \Pr(A \mid K)$ be the government's audit probabilities after each message. The non-compliant firm's expected payoffs are:

$$U_{NC}(D) = \alpha(-f) + (1 - \alpha)g = g - \alpha(f + g)$$

$$U_{NC}(K) = \beta(-f) + (1 - \beta)g = g - \beta(f + g)$$

The non-compliant firm weakly prefers D over K iff:

$$g - \alpha(f + g) \geq g - \beta(f + g) \iff \alpha \leq \beta$$

The non-compliant firm prefers disclosure whenever the government audits after D no more than after K . Intuitively, disclosure is only attractive if it does not attract more scrutiny than concealment.

Posterior under full pooling on D .

If both types play D , observing D is uninformative. By Bayes' rule:

$$\mu(D) = \Pr(NC | D) = \frac{(1-p) \cdot 1}{p \cdot 1 + (1-p) \cdot 1} = 1-p$$

With $p = \frac{2}{3}$: $\mu(D) = \frac{1}{3}$.

Is auditing rational at this posterior?

From part (c), the government strictly prefers to audit iff $\mu > \frac{c}{f} = \frac{1}{4}$. Since $\frac{1}{3} > \frac{1}{4}$, the government strictly prefers to audit after D under the pooling posterior. So $\alpha = 1$.

Does the non-compliant firm strictly prefer to remain at D ?

With $\alpha = 1$, the non-compliant firm earns $-f = -4$ from D . Its payoff from switching to K is $g - \beta(f + g) = 3 - 7\beta$. For the non-compliant firm to weakly prefer D we need:

$$-f \geq g - \beta(f + g) \iff \beta \geq 1$$

Since $\beta \leq 1$ by definition, this holds only at $\beta = 1$. At $\beta = 1$ the non-compliant firm earns $-f$ from both D and K and is therefore **indifferent**, not strictly preferring D . There is no value of β that makes the non-compliant firm strictly prefer D under pooling.

How the answer depends on the government's belief after K .

The information set after K is never reached when both types play D , so Bayes' rule does not pin down $\mu(K)$. The government can assign any belief there.

Conclusion. The non-compliant firm cannot strictly prefer D under pooling. At best, it is indifferent. Whether it strictly prefers to switch to K depends on the belief assigned after K . If

$$\mu(K) \geq \frac{1}{4},$$

then auditing after K is rational and the non-compliant firm is indifferent between D and K . If

$$\mu(K) < \frac{1}{4},$$

then the government does not audit after K , and the non-compliant firm strictly prefers K .

Part (e): The Non-Compliant Firm's Indifference Condition

We ask: given that the compliant type plays D with probability 1 and the non-compliant type plays D with probability $\lambda \in [0, 1]$, what value of λ , together with consistent government audit probabilities α and β , makes the non-compliant firm indifferent between D and K ?

1: Government's best response after K .

Since the compliant type plays D with probability 1:

$$\Pr(K | C) = 0, \quad \Pr(K | NC) = 1 - \lambda$$

For any $\lambda < 1$, observing K identifies the firm as non-compliant with certainty:

$$\mu(K) = \Pr(NC | K) = 1$$

Since $\mu(K) = 1 > \frac{1}{4}$, the government strictly prefers to audit after K :

$$\beta = 1$$

Step 2: Indifference condition for the non-compliant firm.

From part (d), the non-compliant firm is indifferent between D and K iff $\alpha = \beta$. Since $\beta = 1$:

$$\alpha = 1$$

3: Government's posterior after D .

For $\alpha = 1$ to be consistent with the government's posterior, we need $\mu(D) \geq \frac{c}{f}$. The government is exactly indifferent between A and N – and therefore willing to audit – when:

$$\mu(D) = \frac{c}{f}$$

With the compliant type playing D certainly and the non-compliant type playing D with probability λ :

$$\mu(D) = \frac{(1-p)\lambda}{p + (1-p)\lambda}$$

Setting this equal to $\frac{c}{f}$ and solving for λ :

$$\frac{(1-p)\lambda}{p + (1-p)\lambda} = \frac{c}{f}$$

$$f(1-p)\lambda = c[p + (1-p)\lambda]$$

$$(1-p)(f-c)\lambda = cp$$

$$\lambda = \frac{cp}{(1-p)(f-c)}$$

With $p = \frac{2}{3}$, $c = 1$, $f = 4$:

$$\lambda = \frac{1 \cdot \frac{2}{3}}{\frac{1}{3} \cdot 3} = \frac{2}{3}$$

4: Summary of consistent values.

The values consistent with the government's posteriors and the non-compliant firm's indifference are:

$$\alpha = 1, \quad \beta = 1, \quad \lambda = \frac{2}{3}$$

After K , auditing is strictly preferred by the government since $\mu(K) = 1$. After D , auditing is weakly preferred since the posterior is exactly at the cutoff $\mu(D) = \frac{c}{f} = \frac{1}{4}$.

5: Interpretation.

When $\lambda = \frac{2}{3}$, the non-compliant firm plays D with probability $\frac{2}{3}$ and K with probability $\frac{1}{3}$. It is the non-compliant firm that randomizes; the government audits with certainty after both messages. The mixing probability λ is pinned down by the requirement that the government's posterior after D equals $\frac{c}{f}$, which in turn makes the non-compliant firm indifferent between D and K .

The key mechanism: the compliant type's commitment to D means that concealment is a perfect signal of non-compliance ($\mu(K) = 1$), so the government audits after K with certainty. This forces $\beta = 1$, which forces $\alpha = 1$ from the indifference condition, and λ is then back-solved to make the government's posterior consistent with auditing after D .

6: Parameter restriction.

For $\lambda \in (0, 1)$ we need:

$$\frac{cp}{(1-p)(f-c)} < 1 \iff \frac{c}{f} < 1-p$$

An interior value $\lambda \in (0, 1)$ exists if and only if the audit cost-to-fine ratio $\frac{c}{f}$ is strictly less than the prior probability of non-compliance $1-p$. If $\frac{c}{f} \geq 1-p$, the posterior condition $\mu(D) = \frac{c}{f}$ cannot be reached at any interior λ .

With our parameters: $\frac{1}{4} < \frac{1}{3}$. The condition holds.