Players, Strategies, and Payoffs

Game Theory · Module 1 — Strategic Form Games

Node 1.1

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.1

This is the foundation node of the entire game theory course. Every subsequent structural domain — equilibrium, extensive form, information, repetition, mechanism design, evolution, networks, learning, and strategic systems integration — is constructed on top of this single ontological move: the identification of a strategic situation with a triple **(players, strategies, payoffs)**. Before you can solve a game, you must specify what it *is*, and the normal-form definition is the canonical answer. It is the game-theoretic analogue of "Hilbert space + operators + states" in quantum mechanics, or "sample space + sigma algebra + measure" in probability theory: a minimal ontological skeleton onto which everything else hangs.

Formal Definition

A **strategic form game** (or **normal form game**) is a triple

$$G = (N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$$

where:

- $N = \{1, 2, \dots, n\}$ is a finite set of **players**.
- S_i is the set of **pure strategies** available to player i. A pure strategy $s_i \in S_i$ is a complete specification of what player i will do. The **joint strategy space** is the Cartesian product $S = S_1 \times S_2 \times \cdots \times S_n$, and an element $s = (s_1, \dots, s_n) \in S$ is called a **strategy profile**.
- $u_i : S \to \mathbb{R}$ is the **payoff function** of player i. It assigns a real-valued payoff to every strategy profile. The **payoff vector** at profile s is $u(s) = (u_1(s), \dots, u_n(s)) \in \mathbb{R}^n$.

A game is **finite** if N is finite and every S_i is finite. A game is **two-player zero-sum** if $|N| = 2$ and $u_1(s) + u_2(s) = 0$ for all $s \in S$. A game is a **constant-sum** game if this sum is a constant for all profiles.

It is standard notation to write $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$ for the strategy profile of all players *other* than i, so that $s = (s_i, s_{-i})$. Player i's payoff can then be written $u_i(s_i, s_{-i})$, emphasizing the variable the player actually controls.

The entire subsequent theory — dominance, best response, Nash equilibrium, mixed strategies, evolutionary stability, learning dynamics — takes this triple as its input and asks: *what strategy profiles can we justify as outcomes of rational strategic reasoning?*

Intuition

Three ingredients, each answering one question about the strategic situation:

1. **Who is interacting?** $\to$ Players N.
2. **What can each one do?** $\to$ Strategy sets $\{S_i\}$.
3. **What does each one care about?** $\to$ Payoff functions $\{u_i\}$.

The structural move is that payoffs depend on the *entire* profile, not just on a player's own choice. This is the feature that makes game theory *game* theory rather than *decision* theory. In a decision problem, you choose s_i to maximize $u_i(s_i)$ — a one-player optimization. In a game, you choose s_i to maximize $u_i(s_i, s_{-i})$, and s_{-i} is chosen by other agents whose own choices depend on yours. This mutual dependence is the source of both the mathematical difficulty and the conceptual richness of the field.

The payoff function is a full specification of preferences. It already bakes in risk attitudes, time preferences, altruism, and anything else you want to say about how the player values outcomes. The axiomatic foundation (von Neumann–Morgenstern expected utility, 1944) shows that under mild rationality assumptions, any preference ordering over outcome lotteries is representable by a real-valued utility function, which is what u_i is.

One subtle point: a strategy is a **complete plan of action**, not a single move. In a normal-form game this distinction is trivial (each s_i is just a label), but when we lift to extensive form (Module 3), a strategy will be "what I would do at every information set I could reach" — a function from my information states to actions. The normal form is the flattening of that richer object onto a finite-dimensional table.

Structural Role

1.1 is the minimal ontological commitment on which the entire course rests:

- **1.2 Pure vs mixed strategies** extends S_i from pure points to probability distributions $\Delta(S_i)$ over pure points, allowing randomization.
- **1.3 Dominance** uses u_i to compare strategies pointwise across opponent profiles.
- **1.4 Best response** is the function $B_i(s_{-i}) = \arg\max_{s_i \in S_i} u_i(s_i, s_{-i})$, which takes opponent profiles as inputs and returns optimal own-strategies.
- **1.5 Zero-sum and minimax** restricts attention to games where $u_1 + u_2 = 0$, which has special structure (von Neumann's minimax theorem).
- **1.6 Normal form as universal representation** shows that every game, even an extensive-form one with thousands of decision nodes, can be flattened to this triple — possibly at the cost of an enormous strategy space.

From here, the entire equilibrium theory of Module 2 takes off: a **Nash equilibrium** is a profile s^* such that $s_i^* \in B_i(s_{-i}^*)$ for every i, which is a fixed point of the joint best-response correspondence. None of that can even be stated without the definitions of N, S_i, and u_i.

The definition also exports cleanly to every downstream domain. In extensive-form games (Module 3), S_i becomes the set of *plans* — functions from information sets to actions. In Bayesian games (Module 4), it becomes a set of *type-contingent* plans. In mechanism design (Module 6), the designer chooses u_i indirectly through the mechanism. In evolutionary games (Module 7), S_i is a population distribution and u_i is the expected fitness. The triple (N, S, u) is the rigid ontological core that survives every refinement.

Mathematical Development

The payoff matrix

For a two-player finite game, the payoff functions are conveniently displayed as a pair of matrices. Let $S_1 = \{a_1, \ldots, a_m\}$ be Row's strategies and $S_2 = \{b_1, \ldots, b_k\}$ be Column's. Define

$$A_{ij} = u_1(a_i, b_j), \qquad B_{ij} = u_2(a_i, b_j).$$

For a zero-sum game, $B = -A$ and only one matrix is needed. For a general-sum game, the "bimatrix" (A, B) — often written entry by entry as (A_{ij}, B_{ij}) — encodes the whole game.

Example: Matching Pennies. Two players each choose Heads or Tails. Row wins +1 if they match, Column wins +1 if they differ.

$$A = \begin{pmatrix} +1 & -1 \\ -1 & +1 \end{pmatrix}, \qquad B = -A.$$

Example: Prisoner's Dilemma. Two suspects each choose Cooperate (C) or Defect (D). Payoffs (u_1, u_2):

	C	D
C	$(3,3)$	$(0,5)$
D	$(5,0)$	$(1,1)$

Example: Stag Hunt. Two hunters each choose Stag (large cooperative prize) or Hare (small safe prize):

	Stag	Hare
Stag	$(4,4)$	$(0,3)$
Hare	$(3,0)$	$(3,3)$

These three games — Matching Pennies, Prisoner's Dilemma, Stag Hunt — span the strategic archetypes of the entire theory. Matching Pennies is irreducibly mixed (no pure equilibrium). Prisoner's Dilemma has a dominant strategy that produces a Pareto-inferior outcome. Stag Hunt has multiple equilibria with different risk profiles. Learn to recognize which type you're in.

Strategy profile space

The joint strategy space is the Cartesian product

$$S = \prod_{i \in N} S_i,$$

with cardinality $\prod_i |S_i|$. For an n-player game with k strategies each, $|S| = k^n$ — the strategy space grows exponentially in the number of players. This is the combinatorial cost of representing a game in normal form.

Best-response function (preview of 1.4)

$$B_i(s_{-i}) = \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i' \in S_i\}.$$

B_i is set-valued (a **correspondence**) because there may be ties. When the maximizer is unique for every s_{-i}, it is a genuine function. The central equilibrium concept of Module 2 — Nash equilibrium — is a fixed point of the joint correspondence $B(s) = (B_1(s_{-1}), \ldots, B_n(s_{-n}))$.

Utility scaling invariance

Payoff functions are only determined up to positive affine transformations: $u_i' = \alpha_i u_i + \beta_i$ with $\alpha_i > 0$ produces a strategically equivalent game. All solution concepts (dominance, best response, Nash equilibrium) are invariant under such transformations. This means the "units" of payoff are irrelevant — only the *ordering* and *ratios of differences* matter. In particular, you can add a constant to any row of Row's payoff matrix without changing the strategic content.

Worked Example

A concrete three-player game

Three competitors — Alice, Bob, Carol — each choose between Invest (I) or Wait (W). The payoffs, (u_A, u_B, u_C):

- If all three Invest: $(2, 2, 2)$. Market saturated.
- If two Invest, one Waits: investors get 1 each, waiter gets 4. Waiter free-rides.
- If one Invests, two Wait: investor gets -1, waiters get 0. Investor bears the cost alone.
- If all three Wait: $(0, 0, 0)$. Nothing happens.

Players: $N = \{A, B, C\}$. **Strategies:** $S_A = S_B = S_C = \{I, W\}$. **Profile space:** $S = \{I, W\}^3$, with 8 elements. **Payoffs:** $u_A, u_B, u_C : S \to \mathbb{R}$ defined by the list above.

Best response of Alice. Condition on $s_{-A} = (s_B, s_C)$:

- (I, I): Alice gets 2 (I) vs 4 (W) → play W.
- (I, W): Alice gets 1 (I) vs 0 (W) → play I.
- (W, I): Alice gets 1 (I) vs 0 (W) → play I.
- (W, W): Alice gets -1 (I) vs 0 (W) → play W.

So Alice's best-response function is "Invest iff exactly one of Bob/Carol Invests." This already hints at the equilibrium structure: the game has pure equilibria of the form "exactly two invest" (who the two are matters but the waiter free-rides) plus the all-Wait equilibrium where no one moves. Module 2 will make this precise.

Strategic archetype. This game is a three-player generalization of a "volunteer's dilemma" with free-riding. The interdependence is obvious: Alice's best action depends on Bob's and Carol's choices, and symmetrically.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Player set N	Seats in the hand (HU, 6-max, full ring — up to 9 seats)
Strategy set S_i	The complete set of decision rules available — your full strategy, not a single decision
Pure strategy s_i	A fully specified plan: "with AA, 3-bet to 9bb; with 72o, fold; with AK on Q-high flop after 3-betting, c-bet 75% pot; etc."
Mixed strategy σ_i	A randomized plan: "with AJs in MP vs UTG open, call 60% / 3-bet 40%"
Payoff $u_i(s)$	Expected winnings (in chips, BB, or $) of player i given the whole-table strategy profile
Zero-sum structure	Cash poker (minus rake) is zero-sum across the players at the table

The key conceptual move: **a poker strategy is not "what I do with this hand in this spot" — it is the entire decision rule covering every possible hole-card / board / action-history combination.** When a solver outputs "the GTO strategy," it is outputting $s_i \in S_i$ in the game-theoretic sense: a function from every information state to an action distribution.

Concrete situation

Heads-up no-limit hold'em, 100 BB effective. You are the BTN (small blind), opponent is BB. Preflop only. Simplify: each player has either a strong hand (S, top 20% of hands), medium (M, next 30%), or trash (T, bottom 50%). Actions: you can Open (O) or Fold (F). If you Open, BB can Call (C) or Fold (F).

This is a genuine normal-form game once you enumerate types and actions. The strategy set for you is $S_{BTN} = \{O, F\}^3$ (one action per hand class) — **8 pure strategies**, not 2. For example, "Open with S, Open with M, Fold with T" is one pure strategy; "Open with everything" is another; "Fold with everything" is another. The strategy is type-contingent. Similarly BB has $S_{BB} = \{C, F\}^3$ if they commit to call or fold per type (after conditioning on BTN having opened — this is where information sets come in, Module 3).

Payoff. $u_{BTN}(s)$ is the expected chip winnings given the joint strategy, computed over the joint distribution of hand classes and the equity of each class vs opponent's calling range. This is exactly a normal-form payoff function.

What this understanding buys you

Recognizing that your "poker strategy" is the entire conditional plan — not a collection of individual decisions — is the move that unlocks everything. Every discussion of GTO, every solver output, every hand review is working with $s_i \in S_i$ in this sense. When a poker coach says "balance your range," they are saying "construct your s_i so that your action at each information set reveals as little as possible about your type." When they say "your line should be the same with your strong and weak hands in this spot," they are saying your s_i should assign the same action to both types at that information set. The ontology of 1.1 is the ontology of every serious strategic poker discussion.

Cross-Links

- **1.2 Pure vs mixed strategies** — extending S_i to probability distributions.
- **1.4 Best-response correspondences** — the central tool for analyzing strategic situations.
- **2.1 Nash equilibrium** — the fixed-point concept built directly on this triple.
- **3.1 Game trees** — the extensive-form perspective, which flattens back to normal form.
- **Statistics 2.x Decision theory** — single-player special case of this framework.
- **Systems thinking 2.1 State spaces** — S is the state space of the strategic system.
- **Economics** — preferences, choice, utility theory is the single-player foundation.

Structural Compression

Invariant: Every strategic situation reduces to a triple $(N, \{S_i\}, \{u_i\})$ of players, strategy spaces, and payoff functions, with payoffs depending on the entire joint profile.

Mechanism: Cartesian-product construction of the joint strategy space; real-valued payoff functions defined on the joint space; each player controls only their own coordinate.

Cross-System Echo: Compare to the postulates of quantum mechanics (state space + observables + unitary dynamics), the axioms of probability (sample space + sigma-algebra + measure), and the ontology of statistical models (data space + parameter space + likelihood). Each is a minimal ontological skeleton onto which a rich theory is then built. Game theory's skeleton is (N, S, u).

Exercises

1. Write out the payoff matrices (as a bimatrix) for the Battle of the Sexes game: two players choose between Opera and Football; each prefers being together, but Row prefers Opera and Column prefers

Football. Payoff 2 for preferred joint activity, 1 for non-preferred joint activity, 0 for being apart.

2. For the three-player volunteer game in the Worked Example, write out the payoff function u_A as a table indexed by (s_B, s_C) and verify the best-response computation in the text.

3. Show that adding a constant c_i to every payoff of player i (i.e. $u_i \to u_i + c_i$) does not change the set of best responses of any player. Then show that multiplying player i's payoffs by a positive scalar $\alpha_i > 0$ also preserves best responses. Conclude that solution concepts based on best responses are invariant under positive affine transformations.

4. Convert the following poker situation to normal form: two players, each dealt one card from a 3-card deck $\{A, K, Q\}$. Each antes 1 chip. Player 1 acts first and may Bet 1 or Check. If P1 Bets, P2 may Call or Fold. If P1 Checks, the hand goes to showdown. Higher card wins the pot at showdown. Write the strategy sets S_1, S_2 (each player's complete plan as a function of the card they hold) and the payoff function u_1 (average winnings). How many pure strategies does each player have?

5. A finite two-player game has $|S_1| = m$, $|S_2| = k$. How many distinct games are there up to a relabeling of strategies, if payoffs are drawn from $\{0, 1, 2\}$? How does this count grow with m, k?

6. A **dominant strategy** for player i is a strategy s_i^* such that $u_i(s_i^*, s_{-i}) \geq u_i(s_i', s_{-i})$ for all $s_i' \in S_i$ and all $s_{-i} \in S_{-i}$. Show that in the Prisoner's Dilemma, Defect is a dominant strategy for each player. Then show that in the Stag Hunt, *no* strategy is dominant for either player.

7. **Argue, structurally**, why the normal-form triple (N, S, u) is strictly less expressive than an extensive-form description of the same situation. Identify at least one strategic fact (a feature of the situation that affects how rational players should reason) that the normal form loses compared to the extensive form. (Hint: think about the order in which players move, and what they know when they move.)

Game Theory · Module 1 — Strategic Form Games · Node 1.1
Concepts: probability-space
Enables: 1.2, 1.3, 1.4, 1.5, 1.6