Pure vs Mixed Strategies

Game Theory · Module 1 — Strategic Form Games

Node 1.2

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.2

This node enlarges the strategy object. Node 1.1 defined a pure strategy as a deterministic choice from S_i. Many strategic situations — any situation whose opponent would exploit a deterministic rule — require randomization. A *mixed strategy* is a probability distribution over pure strategies, and the lift from S_i to the simplex $\Delta(S_i)$ is the single move that makes Nash equilibrium always exist in finite games (Module 2.4). Without this lift, most interesting games have no solution concept at all. It is also the node where expected utility enters game theory as the working notion of payoff.

Formal Definition

Let $G = (N, \{S_i\}, \{u_i\})$ be a finite strategic form game. A **mixed strategy** for player i is a probability distribution over S_i:

$$\sigma_i \in \Delta(S_i) = \left\{ \sigma_i : S_i \to [0,1] \;\middle|\; \sum_{s_i \in S_i} \sigma_i(s_i) = 1 \right\}.$$

The **support** of σ_i is the set $\operatorname{supp}(\sigma_i) = \{s_i \in S_i : \sigma_i(s_i) > 0\}$. A **pure strategy** is the special case in which σ_i places unit mass on a single point; equivalently, it is a vertex of $\Delta(S_i)$.

A **mixed strategy profile** is $\sigma = (\sigma_1, \ldots, \sigma_n) \in \Delta(S_1) \times \cdots \times \Delta(S_n)$. Under the assumption that players randomize independently, the profile induces a product distribution over $S = \prod_i S_i$:

$$\sigma(s) = \prod_{i \in N} \sigma_i(s_i).$$

Player i's **expected payoff** under the mixed profile is obtained by extending u_i linearly:

$$U_i(\sigma) = \sum_{s \in S} \sigma(s)\, u_i(s) = \sum_{s \in S} \left(\prod_{j \in N} \sigma_j(s_j) \right) u_i(s).$$

The extended payoff $U_i : \prod_j \Delta(S_j) \to \mathbb{R}$ is **multilinear**: linear in each σ_j holding the others fixed. This is the single most important technical fact about mixed strategies.

Intuition

There are three reasons to randomize.

First: to avoid being read. In a zero-sum game with no pure equilibrium — Matching Pennies is the canonical example — any deterministic plan can be exploited by the opponent. Randomization makes your action unpredictable, which is exactly the thing a hidden coin provides. The minimum information you can leak about your action is achieved precisely when you are drawing from a well-chosen distribution.

Second: to make opponents indifferent. In mixed equilibria, your distribution has the subtle property that your opponent is indifferent between the strategies in the support. This is a stability condition: if they were not indifferent, they would concentrate on one strategy, and the game would unravel.

Third: to represent a population. In evolutionary and population games (Module 7), a mixed strategy $\sigma_i(s_i)$ reinterprets as the fraction of the population playing s_i. The "randomization" is not in the head of any individual player; it is in the population distribution.

Multilinearity deserves a moment of its own. Because U_i is linear in each coordinate, an expected-payoff maximizer always has a best response that is pure: you cannot strictly prefer a mixture to every vertex it assigns weight to. This gives the fundamental lemma that we will lean on throughout Module 2: *every best response to σ_{-i} is a mixture whose support lies inside the set of pure best responses.*

Structural Role

Mixed strategies enter every downstream node:

- **1.3 Dominance.** A mixed strategy can strictly dominate a pure strategy that no single pure strategy dominates — so iterated elimination must be stated in terms of mixtures.
- **1.4 Best response.** The best-response correspondence extends from $S_{-i} \to S_i$ to $\Delta(S_{-i}) \to \Delta(S_i)$. The image is always a face of the simplex.
- **1.5 Minimax.** The minimax theorem for zero-sum games requires mixed strategies to hold; it fails in pure strategies.
- **2.3 Mixed-strategy Nash equilibrium.** The central equilibrium concept of the course is a fixed point of the best-response correspondence on $\prod_i \Delta(S_i)$.
- **2.4 Nash existence.** Kakutani's fixed-point theorem requires the domain $\prod_i \Delta(S_i)$ to be compact and convex — which it is, precisely because we allowed mixtures.
- **7 Evolutionary dynamics.** Replicator dynamics live on the mixed-strategy simplex.
- **9 Learning.** No-regret algorithms produce sequences of mixed strategies.

Without the lift from S_i to $\Delta(S_i)$, the entire equilibrium theory of Modules 2 and beyond is either vacuous or ill-defined.

Mathematical Development

The simplex $\Delta(S_i)$

For $|S_i| = m$, $\Delta(S_i)$ is an $(m-1)$-dimensional simplex embedded in $\mathbb{R}^m$:

$$\Delta(S_i) = \{x \in \mathbb{R}^m : x_k \geq 0,\ \textstyle\sum_k x_k = 1\}.$$

This set is **compact** (closed and bounded in $\mathbb{R}^m$) and **convex** (convex combinations of probability distributions are probability distributions). Both properties are essential for existence theorems.

Expected payoff is multilinear

Fix opponent profile σ_{-i}. Then

$$U_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) \cdot U_i(s_i, \sigma_{-i}),$$

where $U_i(s_i, \sigma_{-i}) = \sum_{s_{-i}} \sigma_{-i}(s_{-i}) u_i(s_i, s_{-i})$ is the expected payoff of the pure strategy s_i against the mixed profile σ_{-i}. This expression is linear in σ_i — each $\sigma_i(s_i)$ appears to first order. Consequently:

Best response to a mixed profile

$$\mathrm{BR}_i(\sigma_{-i}) = \arg \max_{\sigma_i \in \Delta(S_i)} U_i(\sigma_i, \sigma_{-i}).$$

By linearity, any maximizer is supported on the set $\arg\max_{s_i \in S_i} U_i(s_i, \sigma_{-i})$. If this argmax is a single pure strategy, it is the unique best response. If it contains several pure strategies, any mixture over them is also a best response. Conversely, no mixed strategy assigning positive weight to a non-best-response pure strategy can be a best response.

The indifference condition

If σ_i^* is a best response and $\mathrm{supp}(\sigma_i^*)$ contains two or more pure strategies, then all pure strategies in the support must yield the same expected payoff against σ_{-i}. This is the **indifference principle** and is the technical backbone of every mixed-equilibrium calculation.

Example: Matching Pennies

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

Let Column play $(q, 1-q)$. Row's expected payoff from Heads is $q - (1-q) = 2q - 1$; from Tails is $-q + (1-q) = 1 - 2q$. Row is indifferent iff $2q - 1 = 1 - 2q$, i.e. $q = 1/2$. By symmetry Row must play $(1/2, 1/2)$ in equilibrium. This is the unique mixed equilibrium and there is no pure one.

Example: Rock–Paper–Scissors

With payoff $+1$ win, -1 loss, 0 tie, the unique equilibrium mix is $(1/3, 1/3, 1/3)$ for both players, by a symmetric indifference argument.

Kuhn's purification theorem (preview)

Harsanyi (1973) showed that mixed equilibria of a game G can be recovered as limits of *pure* equilibria of a slightly perturbed game with small private noise in the payoffs. So "randomization" is not a philosophical commitment — it is the strategic residue of tiny type uncertainty. Module 4 develops this fully.

Worked Example

Attacker–Defender

A defender chooses which of two targets, $\{T_1, T_2\}$, to protect; an attacker simultaneously chooses one to attack. Target 1 has value $v_1 = 9$, target 2 has value $v_2 = 4$. If the attacker hits an unprotected target, the defender loses its value and the attacker gains it. If the attacker hits a protected target, payoffs are zero. (Zero-sum.)

Payoff matrix for the defender (rows = Defender, columns = Attacker):

$$A = \begin{pmatrix} 0 & -9 \\ -4 & 0 \end{pmatrix}.$$

Let the defender play $(p, 1-p)$ (probability of protecting T_1, then T_2). The attacker's expected gain from attacking T_1 is $9(1-p)$; from attacking T_2 is $4p$. The attacker is indifferent iff

$$9(1-p) = 4p \quad \Longleftrightarrow \quad p = \frac{9}{13} \approx 0.69.$$

Now let the attacker play $(q, 1-q)$. The defender's expected loss from protecting T_1 is $4(1-q)$; from protecting T_2 is $9q$. Indifference gives $4(1-q) = 9q \Rightarrow q = 4/13 \approx 0.31$. The unique mixed equilibrium is

$$\sigma_D^* = (9/13, 4/13), \qquad \sigma_A^* = (4/13, 9/13).$$

Notice the inversion: the defender protects the *more valuable* target with probability 9/13, and the attacker attacks the more valuable target with probability 4/13 — less than half. The high-value target is so well-defended that the attacker is better off going for the soft one more than half the time. This non-obvious inversion is the content of the mixed equilibrium.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Pure strategy s_i	A 100/0 decision: "always 3-bet this hand in this spot"
Mixed strategy σ_i	A frequency: "3-bet this hand 40%, call 60%"
Support $\operatorname{supp}(\sigma_i)$	The set of actions you ever take with this exact hand + spot
Indifference condition	Solver balance: all actions in your mix must have the same EV
Multilinearity of U_i	EV of a mix is the average EV of its pure parts — you cannot mix into something better than your best pure action
Purification	Noise in chip values / rake rounding makes "mixed" play look like tiny private-value tilts

The table reveals a central fact: **every frequency a solver outputs is an indifference frequency.** If the solver tells you to 3-bet AJs 40% and call 60%, it is because at the current game state both 3-betting and calling yield the same EV — you are indifferent, and the 40/60 is chosen to make your *opponent* indifferent at some downstream node.

Concrete situation

BTN vs BB, 100 BB, single-raised pot. Flop comes $K\heartsuit 7\clubsuit 2\diamondsuit$ — a dry, BTN-range-favoring board. You (BTN) are considering your c-bet strategy with $A\clubsuit 5\clubsuit$.

A pure strategy is "always bet" or "always check." A mixed strategy is "bet x, check $1-x$." A modern solver output for this spot is approximately: bet small (25–33% pot) with $x \approx 0.75$, check with $1 - x \approx 0.25$. Why a mix? Because with $A5s$, the EV of betting small and the EV of checking back are within a fraction of a BB of each other in equilibrium. The solver assigns exactly the weights that make the **opponent's** best-response correspondence balanced: BB's continue vs fold frequencies, their check-raise frequencies, and their turn play all line up only at this exact mix. Change the frequency, and one of BB's strategies becomes strictly better, and the whole equilibrium collapses.

Concretely: if you always bet with $A5s$, BB can exploit by calling more often with mid-pair holdings because you are bluffing too thin. If you always check, your check-back range becomes too strong, and BB can't profitably float. Mixing is how you keep both of BB's counter-strategies honest.

What this understanding buys you

Two immediate operational lessons.

First, you never need to beat yourself up about "choosing the wrong action" when you were in a mixed spot. By construction, the two actions have the same EV. If the solver says 50/50, then whichever one you took is exactly as good as the other — the only mistake is going 100/0 when the equilibrium calls for a mix.

Second, the practical way to "play a mixed strategy" at the table is to use an external randomizer (the second hand of a watch, the suit of your lowest hole card, an app). Humans are notoriously bad at internally generated randomness — we overproduce alternation. Any pattern in your "random" behavior is exactly what good opponents are reading. The indifference principle is fragile: it only works if your randomization is actually random.

Cross-Links

- **1.1 Players, strategies, payoffs** — the base definition this node extends.
- **1.3 Dominance** — iterated elimination must be stated for mixed strategies.
- **1.4 Best-response correspondences** — now taking values in $\Delta(S_i)$.
- **1.5 Zero-sum and minimax** — minimax only holds after the mixed-strategy lift.
- **2.3 Mixed-strategy Nash equilibrium** — the solution concept this node enables.
- **2.4 Nash existence theorem** — requires $\Delta(S_i)$ compact convex.
- **4 Information and beliefs** — Harsanyi purification relates mixing to type uncertainty.
- **7 Evolutionary dynamics** — mixed strategies as population distributions.
- **Statistics — expected value and linearity of expectation** — the purely technical foundation of U_i.

Structural Compression

Invariant: The strategy set of each player lifts from the finite set S_i to the simplex $\Delta(S_i)$, and the payoff function u_i lifts to its multilinear extension U_i. Pure strategies are the vertices of the simplex; mixed strategies are its interior points and faces.

Mechanism: Independent randomization on each coordinate produces a product distribution on the joint strategy space, and multilinearity of U_i gives an indifference characterization of best responses.

Cross-System Echo: The move from S_i to $\Delta(S_i)$ is the same move as from Boolean logic to probability, from classical states to quantum mixed states (density matrices), and from deterministic MDPs to stochastic policies. In each case, convexification of the decision object rescues existence of a solution.

Exercises

1. Verify that in Matching Pennies the unique Nash equilibrium is $((1/2, 1/2), (1/2, 1/2))$, and that no pure-strategy equilibrium exists. Compute each player's expected payoff at the equilibrium.

2. In the Attacker–Defender game from the worked example, change $v_1 = 16, v_2 = 4$. Recompute the equilibrium. Confirm that the defender still protects the more valuable target with probability strictly greater than $1/2$, and that the attacker's attack probability on the high-value target goes *down* as v_1 increases.
3. Prove the indifference principle: if σ_i^* is a best response to σ_{-i} and $s_i, s_i' \in \text{supp}(\sigma_i^*)$, then $U_i(s_i, \sigma_{-i}) = U_i(s_i', \sigma_{-i})$.
4. Show that the expected payoff function $U_i(\sigma_i, \sigma_{-i})$ is continuous in (σ_i, σ_{-i}) as a function on $\prod_j \Delta(S_j)$. (This is used in the Nash existence proof.)
5. Consider a 3-strategy Rock–Paper–Scissors variant where Rock beats Scissors by a margin of 2, and the other wins are by margin 1 (standard). Find the unique mixed equilibrium. Is it still $(1/3, 1/3, 1/3)$? If not, explain what changes.
6. A player's strategy is called **completely mixed** if every pure strategy is in the support. Show that in a completely mixed equilibrium, every pure strategy in S_i gives the same expected payoff against σ_{-i}^*. Conclude that completely mixed equilibria satisfy $|S_i| - 1$ linear equations per player.
7. **Argue, structurally,** why the lift from S_i to $\Delta(S_i)$ is unavoidable if we want a general existence theorem for equilibrium in finite games. Specifically: give a finite game with no pure Nash equilibrium and explain exactly which step of the would-be existence proof fails in pure strategies but succeeds after the lift. (Hint: think about compactness and convexity of the domain of a fixed-point map.)

Game Theory · Module 1 — Strategic Form Games · Node 1.2
Concepts: probability-space, expectation-value, convexity
Prerequisites: 1.1
Enables: 1.3, 1.4, 1.5, 2.1, 2.3

hbar.university — own.your.cognition