Mixed Strategies

IE 5545: Decision Analysis
Prof. Krishnamurthy Iyer*

May 2, 2026

Until now, we made the assumption that each player chooses a strategy deterministically in the game. However, many instances in practice involve players choosing strategies at random. For example, in rock-paper-scissors, a player may, reasonably, decide to choose one among her three choices at random — possibly to throw her opponent off. Such random choices cannot be described in our current framework. Thus, we must extend our framework to accommodate such random behavior on the part of the players.

As a motivation for settings where random choice of strategies might make sense in the analysis of a game, consider the following example:

	L	M	R
U	$(2,3)$	$(3,0)$	$(2,1)$
D	$(3,0)$	$(4,3)$	$(1,1)$

It is straightforward to see that this game has no strictly dominated or weakly dominated strategies. Thus, the ISD process stops without eliminating any strategy. In other words, this game has no strict DSE, weak DSE, or ISD-equilibrium. Thus, it seems that we are helpless if we restrict ourselves to using dominance concepts.

However, note that if we allow player 2 to choose a strategy randomly, then consider a "strategy" for player 2 where she chooses L with probability 0.5 and M with probability 0.5. Call this "strategy" λ. Appending this "strategy" to the payoff matrix yields the following:

	L	M	R	λ
U	$(2,3)$	$(3,0)$	$(2,1)$	$(2.5,1.5)$
D	$(3,0)$	$(4,3)$	$(1,1)$	$(3.5,1.5)$

Here, the payoffs under λ correspond to the average of the payoffs under L and M (for each strategy of player 1).

From this matrix, we see that, irrespective of how P1 plays, P2 gets strictly higher *expected payoff* on playing λ than on playing R. Thus, we can say that R is strictly dominated by the random "strategy" λ. If we eliminate R, then in the reduced game, by ISD, we obtain (D, M) as the solution. (Exercise: verify).

The above example illustrates that we can make use of such random strategies, referred to as a "mixed strategies", to extend our notion of dominance to obtain solutions to larger classes of games. This motivates the study of mixed strategies. We will now formally extend the notion of a strategy (and the payoffs) to include such randomizations.

1 Formal definitions

Let the set of (deterministic) strategies for player i be

$$\mathcal{S}_i = \{s_i^1, s_i^2, \ldots, s_i^k\}.$$

We refer to these deterministic strategies as "pure", meaning they involve no randomization.

A *mixed strategy* σ_i is defined as a probability distribution over $\mathcal{S}_i$. For each $s_i \in \mathcal{S}_i$, we let $\sigma_i(s_i)$ denote the probability with which a particular pure strategy s_i is played under σ_i. We will often denote a mixed strategy using the column vector notation:

$$\sigma_i = \begin{bmatrix} p_i^1 \\ p_i^2 \\ \vdots \\ p_i^k \end{bmatrix} = [p_i^1 \;\; p_i^2 \;\; \cdots \, p_i^k]^\top,$$

where $p_i^j = \sigma_i(s_i^j)$ for all j. Note that we must have $\sum_{j=1}^k p_i^j = 1$ and $p_i^j \geq 0$ for all j.

We let Σ_i denote the set of all mixed strategies of player i. By definition, this is the set of all distributions over $\mathcal{S}_i$. In other words, we have

$$\Sigma_i = \Delta(\mathcal{S}_i) = \left\{ [p_i^1 \;\; p_i^2 \;\; \cdots \, p_i^k]^\top : \sum_{j=1}^k p_i^j = 1, \quad p_i^j \geq 0 \right\}.$$

Note that by setting $p_i^j = 1$ and $p_i^\ell = 0$ for all $\ell \neq j$, we obtain a mixed strategy σ_i that chooses the pure strategy s_i^j with probability 1; in other words playing this mixed strategy σ_i is same as playing the pure strategy s_i^j. For this reason, we note that any pure strategy is also a mixed strategy. Thus, the set of mixed strategies Σ_i can be viewed as a superset of the set of pure strategies $\mathcal{S}_i$ (and hence the mixed strategies are an *extension* of pure strategies).

Example. Suppose $\mathcal{S}_1 = \{U, D\}$, and $\mathcal{S}_2 = \{L, M, R\}$. Consider a mixed strategy σ_1 for player 1 where she chooses between U and D uniformly at random. Further, consider a mixed strategy σ_2 for player 2 where she chooses L with probability 0.4, M with probability 0.2 and R with probability 0.4. Then we have

$$\begin{aligned} &\sigma_1(U) = 0.5 \qquad &&\sigma_1(D) = 0.5 \\ &\sigma_2(L) = 0.4 \qquad &&\sigma_2(M) = 0.2 \qquad &&\sigma_2(R) = 0.4. \end{aligned}$$

Equivalently, in vector notation, we have

$$\sigma_1 = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} \qquad \sigma_2 = \begin{bmatrix} 0.4 \\ 0.2 \\ 0.4 \end{bmatrix} \qquad \square$$

A *mixed strategy profile* is a collection of mixed strategies, one for each player. In a mixed strategy profile $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_N)$, player 1 plays mixed strategy σ_1, player 2 plays mixed strategy σ_2, and so on. Similarly, we let $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_N)$ denote the mixed strategy profile of all players except player i.

Under a mixed strategy profile σ, we assume that each player i is choosing a (pure) strategy s_i randomly but *independently*. In other words, the players' do not correlate their randomization; each is using an independent random number generator. Because the randomizations are independent, we can compute the probability that a pure strategy $s = (s_1, \ldots, s_N)$ is played in the game using just the individual marginal probabilities. In particular, the probability that a pure strategy profile $s = (s_1, \ldots, s_N)$ is played in the game is equal to the product (over all players i) of the probabilities that the pure strategy s_i is played by player i.

$$
\begin{aligned}
&\mathbf{P}\big(s = (s_1, \ldots, s_N) \text{ is realized under } \sigma\big) \\
&= \mathbf{P}\big(s_1 \text{ is realized under } \sigma_1\big) \cdot P\big(s_2 \text{ is realized under } \sigma_2\big) \cdots P\big(s_N \text{ is realized under } \sigma_N\big) \\
&= \sigma_1(s_1) \cdot \sigma_2(s_2) \cdots \sigma_N(s_N).
\end{aligned}
$$

Using this distribution over the pure strategy profiles, we can extend the notion of the payoff function to mixed strategy profiles. The payoff $\pi_i(\sigma)$ to player i when strategy profile $\sigma = (\sigma_1, \ldots, \sigma_N)$ is played is the resulting expected payoff when the (pure) strategy profile is distributed according to the distribution above:

$$
\begin{aligned}
\pi_i(\sigma) &= \sum_{\substack{\text{all pure} \\ \text{strategy profiles } s}} \mathbf{P}\big(s \text{ is realized under } \sigma\big) \cdot \pi_i(s) \\
&= \sum_{s=(s_1,s_2,\ldots,s_N)} \big[\sigma_1(s_1)\,\sigma_2(s_2) \cdots \sigma_N(s_N)\big] \cdot \pi_i(s_1, s_2, \ldots, s_N) \\
&= \sum_{s_1 \in \mathcal{S}_1} \sum_{s_2 \in \mathcal{S}_2} \cdots \sum_{s_N \in \mathcal{S}_N} \sigma_1(s_1)\,\sigma_2(s_2) \cdots \sigma_N(s_N) \cdot \pi_i(s_1, s_2, \ldots, s_N)
\end{aligned}
$$

Example. Consider again the two player game below:

	L	M	R
U	$(2,3)$	$(3,0)$	$(2,1)$
D	$(3,0)$	$(4,3)$	$(1,1)$

Consider the mixed strategy profile $\sigma = (\sigma_1, \sigma_2)$ where

$$
\begin{aligned}
&\sigma_1(U) = 0.5 \qquad && \sigma_1(D) = 0.5 \\
&\sigma_2(L) = 0.4 \qquad && \sigma_2(M) = 0.2 \qquad && \sigma_2(R) = 0.4.
\end{aligned}
$$

Under this profile, the probability that each pure strategy profile is played is given by

$$
\mathbf{P}((s_1, s_2) \text{ is played under } \sigma) = \begin{cases}
0.5 \cdot 0.4 = 0.2 & \text{if } (s_1, s_2) = (U, L); \\
0.5 \cdot 0.2 = 0.1 & \text{if } (s_1, s_2) = (U, M); \\
0.5 \cdot 0.4 = 0.2 & \text{if } (s_1, s_2) = (U, R); \\
0.5 \cdot 0.4 = 0.2 & \text{if } (s_1, s_2) = (D, L); \\
0.5 \cdot 0.2 = 0.1 & \text{if } (s_1, s_2) = (D, M); \\
0.5 \cdot 0.4 = 0.2 & \text{if } (s_1, s_2) = (D, R);
\end{cases}
$$

Then, the payoff of the two played under σ is given by

$$\begin{aligned}
\pi_1(\sigma_1, \sigma_2) &= \sum_{(s_1,s_2)} \mathbf{P}((s_1, s_2) \text{ is played under } \sigma)\pi_1(s_1, s_2) \\
&= 0.2 \cdot 2 + 0.1 \cdot 3 + 0.2 \cdot 2 + 0.2 \cdot 3 + 0.1 \cdot 4 + 0.2 \cdot 1 \\
&= 2.3 \\
\pi_1(\sigma_1, \sigma_2) &= \sum_{(s_1,s_2)} \mathbf{P}((s_1, s_2) \text{ is played under } \sigma)\pi_2(s_1, s_2) \\
&= 0.2 \cdot 3 + 0.1 \cdot 0 + 0.2 \cdot 1 + 0.2 \cdot 0 + 0.1 \cdot 3 + 0.2 \cdot 1 \\
&= 1.3. \qquad \square
\end{aligned}$$

For two player games, the matrix notation provides a convenient way to represent the payoffs. Consider a two player game with m pure strategies for player 1 and n pure strategies for player 2. Recall that we can represent the payoffs of the two players as two matrices Π_1 and Π_2, where Π_i is a $m \times n$ matrix with the $(k, \ell)^{th}$ entry denoting player i's payoff when player 1 plays pure strategy k and player 2 plays pure strategy ℓ.

A mixed strategy σ_1 for player 1 can be represented as an $m \times 1$ column vector. Similarly, a mixed strategy σ_2 for player 2 can be represented as an $n \times 1$ column vector. With this representations, we can write the payoff of player i under a mixed strategy profile $\sigma = (\sigma_1, \sigma_2)$ as

$$\pi_i(\sigma_1, \sigma_2) = \sum_{s_1 \in \mathcal{S}_1} \sum_{s_2 \in \mathcal{S}_2} \sigma_1(s_1)\, \sigma_2(s_2)\, \pi_i(s_1, s_2) = \sigma_1^{\mathsf{T}} \Pi_i \sigma_2.$$

Example. (contd) Using the vector notation, the payoff matrix are given by

$$\Pi_1 = \begin{bmatrix} 2 & 3 & 2 \\ 3 & 4 & 1 \end{bmatrix}, \qquad \Pi_2 = \begin{bmatrix} 3 & 0 & 1 \\ 0 & 3 & 1 \end{bmatrix}.$$

The mixed strategies σ_1 and σ_2 in vector notation are given by

$$\sigma_1 = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} \quad \text{and} \quad \sigma_2 = \begin{bmatrix} 0.4 \\ 0.2 \\ 0.4 \end{bmatrix}.$$

Then, the payoff of the two players under $\sigma = (\sigma_1, \sigma_2)$ can be computed as

$$\begin{aligned}
\pi_1(\sigma_1, \sigma_2) = \sigma_1^T \Pi_1 \sigma_2 = \begin{bmatrix} 0.5 & 0.5 \end{bmatrix} \begin{bmatrix} 2 & 3 & 2 \\ 3 & 4 & 1 \end{bmatrix} \begin{bmatrix} 0.4 \\ 0.2 \\ 0.4 \end{bmatrix} = 2.3 \\
\pi_2(\sigma_1, \sigma_2) = \sigma_1^T \Pi_2 \sigma_2 = \begin{bmatrix} 0.5 & 0.5 \end{bmatrix} \begin{bmatrix} 3 & 0 & 1 \\ 0 & 3 & 1 \end{bmatrix} \begin{bmatrix} 0.4 \\ 0.2 \\ 0.4 \end{bmatrix} = 1.3. \qquad \square
\end{aligned}$$

To summarize, with these definitions, we have extended the notion of a strategic form to include mixed strategies for each player and corresponding payoff functions over mixed strategies. We can analogously define the notions of strict and weak dominance (and ISD) for this extended game. For example, we can now say a pure strategy s_i for player i is strictly dominated by a mixed strategy σ_i, if the payoff of player i under s_i is strictly less than that under σ_i, for any strategy choices of other players. Similarly, we can extend the notion of a Nash equilibrium to mixed strategies; this will be the topic we will take up in detail next.