

Computational Game Theory

Textbook

Tomáš Kroupa

Contents

1	Strategic games	5
1.1	Introduction and motivation	5
1.1.1	From optimization to strategic reasoning	5
1.1.2	From human players to artificial agents	5
1.1.3	Taxonomy of games	6
1.2	Normal-form games and examples	7
1.2.1	Strategy profiles	7
1.2.2	Examples of two-player games	7
1.3	Dominance	9
1.3.1	Rational choice	9
1.3.2	Dominance in the prisoner's dilemma	9
1.3.3	Iterated elimination	9
1.4	Nash equilibrium	10
1.4.1	Best responses	10
1.4.2	Finding pure equilibria	10
1.5	Equilibrium and efficiency	11
1.6	Dynamic games and their normal form	12
1.6.1	The Alice–Bob game	12
1.6.2	Strategies as complete contingent plans	13
1.6.3	Constructing the normal form	13
1.6.4	Nash equilibrium and sequential reasoning	14
2	Mixed-strategy Nash equilibrium	15
2.1	Why randomize?	15
2.2	Mixed strategies and expected utility	15
2.2.1	Probability distributions over pure strategies	15
2.2.2	Expected utility	16
2.3	Nash equilibrium and best responses	17
2.3.1	The best-response correspondence	17
2.3.2	Why pure strategies suffice to test Nash equilibrium	17
2.4	Existence of equilibrium	18
2.5	Supports and indifference	18
2.5.1	The indifference principle	18
2.5.2	The support characterization	19
2.6	Two-player (2×2) examples	20
2.7	Good Samaritan: Who will provide aid?	23
2.7.1	Pure equilibria	23
2.7.2	A symmetric mixed equilibrium	23
2.7.3	The chance that nobody helps	23
	Appendices	25

Normal-form game catalogue	27
Instigation game	27
Prisoner's dilemma	27
Matching pennies	27
Sharing household chores	27
Bach or Stravinski	27
Coordination game	27
Iterated-dominance example	27
Strategic form of the dynamic game	28
Expected utility and partial support	28
Indifference is not sufficient	28
Penalty kicks	28
Professor's dilemma	28
Good Samaritan: two bystanders	28
Extensive-form game catalogue	29
Alice–Bob game	29
References	31

1. Strategic games

1.1 Introduction and motivation

Game theory is the mathematical study of interactive decision-making, where multiple agents pursue their own objectives and each player's outcome depends on the decisions made by all players.

The defining feature is *strategic interdependence*: the outcome depends on the choices of several decision-makers, and what is best for one player can depend on what the others choose. Players may be people, firms, institutions, or autonomous agents. Applications include bargaining, markets, auctions, voting, security, cost allocation, and multi-agent learning.

A utility function represents a player's preferences over outcomes. Utility need not be money, and utility maximization need not mean selfish behavior. If a player values fairness, safety, or another person's welfare, those concerns can be part of that player's preferences.

The modern theory grew from the minimax analysis of two-player zero-sum games [1]. Nash equilibrium supplied a general consistency condition for non-cooperative games [2], while the Shapley value formalized a method for allocating the gains from cooperation [3]. Auction and mechanism design later made strategic reasoning part of the design of markets and institutions [4], [5].

1.1.1 From optimization to strategic reasoning

The standard unilateral optimization treats the environment as fixed. A decision-maker controls the relevant variables and evaluates a single objective. In a game, some variables are controlled by other decision-makers who are optimizing too.

Optimization	Game theory
One decision-maker	Several decision-makers
One objective	Each player has an objective
The decision-maker controls the variables	The outcome depends on everyone's choices
What is my best feasible decision?	What is best when others choose strategically too?

Game theory extends optimization rather than replacing it. Once the opponents' choices are fixed, a player's problem is again an optimization problem. The difficulty is to make the players' optimization problems mutually consistent.

1.1.2 From human players to artificial agents

The word *player* first suggests a person, but the formal concept is more abstract: a player is any decision-making entity to which we assign available strategies and preferences over outcomes. It may therefore represent a person, a firm, an institution, a software agent, or an entire platform.

Strategic interaction among humans has not disappeared, but more decisions are now delegated to algorithms or mediated by computational systems:

- trading algorithms interact in financial markets;
- ride-hailing and airline platforms update prices dynamically;
- recommender systems influence which alternatives users encounter;
- users, developers, and other agents can strategically influence the behavior of systems built around large language models;
- self-play and adversarial training place learned agents in deliberate competition.

Games have also become important testbeds for artificial intelligence: AlphaGo, DeepStack, and AlphaStar demonstrated strong performance in settings with very different action spaces and information structures [6], [7], [8].

This shift raises a conceptual and practical challenge. For an artificial player, we must decide what counts as its information, objective, and strategy; we must also ask whether behavior observed in a test predicts behavior after deployment. Kovařík and coauthors argue that advanced systems may recognize and respond strategically to evaluation itself. Game-theoretic models can make the resulting assumptions explicit and help scrutinize evaluation-based safety claims [9]. Thus games are not only environments in which AI systems act: they can also become controlled tests of competition, coordination, manipulation, and response to incentives.

1.1.3 Taxonomy of games

Games can be classified along several related dimensions. The categories in different rows can be combined: a single game may be dynamic, finite, general-sum, non-cooperative, complete-information, and perfect-information at the same time.

Modelling question	Main alternatives
How are choices represented?	normal form, extensive form
Is the focus on individual choices or coalitions?	non-cooperative, cooperative, coalitional
When are choices made?	static, dynamic
What is observed during play?	perfect information, imperfect information
What is known about the game's structure?	complete information, incomplete information
How large are the strategy sets?	finite, infinite
How are players' interests related?	zero-sum, general-sum

The *normal form* records strategies and utilities while suppressing the order of moves. The *extensive form* represents possible histories, who moves after each history, and what that player observes. A *coalitional game* instead emphasizes what groups of players can achieve together and how their joint value or cost might be allocated.

The information distinctions answer different questions. *Perfect information* means that every player who moves observes the complete history of earlier actions. *Complete information* means that the structure of the game - including the players, strategy sets, and utility functions - is common knowledge. A simultaneous auction can therefore have imperfect information during play and, when valuations are private, incomplete information about the game's payoff-relevant data.

Three examples indicate the scope of the taxonomy:

- A sealed-bid auction is naturally described by simultaneous bids and can be represented in normal form.
- A sequence of observed decisions is naturally drawn as an extensive-form tree.
- Connecting several cities to a shared provider motivates a coalitional cost game: the central question is what each coalition can build and how the cost should be shared.

This chapter develops the first two solution methods for normal-form games: dominance and pure-strategy Nash equilibrium. It also distinguishes strategic stability from social efficiency and shows how a dynamic game can be converted to normal form.

1.2 Normal-form games and examples

Definition

A **normal-form game**, also called a **strategic game**, is a tuple

$$(N, (S_i)_{i \in N}, (u_i)_{i \in N}), \quad (1.1)$$

where

- $N = \{1, \dots, n\}$ is a nonempty player set;
- S_i is the nonempty set of *pure strategies* available to player i ;
- u_i is player i 's utility function,

$$u_i : \prod_{j \in N} S_j \rightarrow \mathbb{R}. \quad (1.2)$$

The game is *finite* when every S_i is finite.

1.2.1 Strategy profiles

The set of pure-strategy profiles is

$$\mathbf{S} = \prod_{i \in N} S_i. \quad (1.3)$$

For player i , write

$$\mathbf{S}_{-i} = \prod_{j \neq i} S_j, \quad \mathbf{s} = (s_i, \mathbf{s}_{-i}). \quad (1.4)$$

Thus $s_i \in S_i$ denotes one strategy of player i , $\mathbf{s} \in \mathbf{S}$ selects one strategy for every player, and $\mathbf{s}_{-i} \in \mathbf{S}_{-i}$ collects the opponents' strategies. This notation makes a unilateral deviation transparent for player i : replace s_i while holding $\mathbf{s}_{-i}$ fixed.

To use a normal-form model, players select their strategies before observing one another's strategy choices:

1. each player i selects one strategy s_i ;
2. the choices form a profile $\mathbf{s} = (s_1, \dots, s_n) \in \mathbf{S}$;
3. player i receives utility $u_i(\mathbf{s})$.

In a dynamic game, an **action** is a move at a particular decision point. A strategy in such a game is a complete contingent plan since it may prescribe later actions after observed histories. In a one-stage game, a pure strategy consists of one action, so the terms are often used interchangeably. Section 1.6 makes this distinction precise.

1.2.2 Examples of two-player games

For a finite two-player game, a payoff matrix places player 1's strategies in rows and player 2's strategies in columns. The cell at (s_1, s_2) contains $(u_1(s_1, s_2), u_2(s_1, s_2))$.

Example: Prisoner's dilemma

Two prisoners independently choose whether to remain quiet (*Cooperate*) or confess (*Defect*). Utilities are the negatives of their prison sentences in years.

	<i>C</i>	<i>D</i>	
<i>C</i>	-1, -1	-4, 0	(1.5)
<i>D</i>	0, -4	-3, -3	

If both remain quiet, each receives a one-year sentence. If only one confesses, that prisoner is released and the other receives four years. If both confess, each receives three years.

A two-player game is **zero-sum** if

$$u_1(s_1, s_2) + u_2(s_1, s_2) = 0 \quad \text{for every } (s_1, s_2) \in S_1 \times S_2. \quad (1.6)$$

Example: Matching pennies

Two players each place a coin under their hand, choosing whether to show heads (*H*) or tails (*T*). They then reveal their coins simultaneously, so neither can observe the other's choice before making their own. If the coins show the same side, player 2 pays one unit to player 1. If they show different sides, player 1 pays one unit to player 2. Taking each player's net gain as their payoff gives the following game:

	<i>H</i>	<i>T</i>	
<i>H</i>	1, -1	-1, 1	(1.7)
<i>T</i>	-1, 1	1, -1	

Player 1 therefore wants the two choices to match, while player 2 wants them to differ. Their interests are exactly opposed: every gain for one player is an equal loss for the other. Neither heads nor tails is always the better choice.

General-sum games need not have either identical or exactly opposed interests.

Example: Sharing household chores

Alice and Bob share a flat, and each has time for one chore before the evening. Alice chooses whether to cook or clean, while Bob chooses whether to shop or take out the trash. They choose independently, without knowing what the other will do.

Alice likes their efforts to contribute to the same goal: her favourite outcome is cooking a meal while Bob shops for groceries, followed by cleaning the flat while he takes out the trash. Bob would rather cover different household needs. His favourite outcome is to restock the groceries while Alice cleans; his second choice is to take out the trash while she cooks. Each assigns a score of 2 to their favourite outcome, 1 to their second choice, and 0 to the remaining two outcomes. Alice's payoff is listed first:

	Shop	Trash	
Cook	2, 0	0, 1	(1.8)
Clean	0, 2	1, 0	

These payoffs represent satisfaction with the combination of chores, rather than payments between the roommates. Their preferences differ, but a gain for one need not be an equal loss for the other.

1.3 Dominance

1.3.1 Rational choice

We model a *rational* player as choosing a strategy that maximizes their utility, given the available information about the opponents' choices. Rationality is conditional: the best strategy can depend on what others do. It also imposes no particular moral content because the utility function itself represents what the player values.

Dominance provides a useful exception to this dependence. If one strategy is worse than another against every possible profile of opponents, a utility-maximizing player can discard it without predicting what the opponents will do.

Definition

For $s_i, t_i \in S_i$:

- t_i is **strictly dominated** by s_i if

$$u_i(s_i, \mathbf{s}_{-i}) > u_i(t_i, \mathbf{s}_{-i}) \quad \text{for every } \mathbf{s}_{-i} \in \mathbf{S}_{-i}; \quad (1.9)$$

- t_i is **weakly dominated** by s_i if

$$u_i(s_i, \mathbf{s}_{-i}) \geq u_i(t_i, \mathbf{s}_{-i}) \quad \text{for every } \mathbf{s}_{-i} \in \mathbf{S}_{-i}, \quad (1.10)$$

with strict inequality for at least one $\mathbf{s}_{-i} \in \mathbf{S}_{-i}$.

A strategy is **strictly dominant** if it strictly dominates every other strategy of that player. Weakly dominant strategies are defined analogously.

A rational player never uses a strictly dominated pure strategy. Eliminating a weakly dominated strategy requires more care because an equality can matter for subsequent strategic comparisons.

1.3.2 Dominance in the prisoner's dilemma

For player 1,

$$u_1(D, C) = 0 > -1 = u_1(C, C), \quad u_1(D, D) = -3 > -4 = u_1(C, D). \quad (1.11)$$

Thus D strictly dominates C . The game is symmetric, so D is also strictly dominant for player 2. The solution based on dominance is therefore (D, D) , without any belief about the other prisoner's choice.

1.3.3 Iterated elimination

Iterated elimination of strictly dominated strategies applies the same reasoning repeatedly:

1. begin with the full strategy sets;
2. find a pure strategy that is strictly dominated in the current restricted game;
3. delete it and form the reduced game;
4. repeat until no further deletion is possible for any player.

Consider the following game between Alice (row player) and Bob (column player) as an example.

$$\begin{array}{cc|cc} & c & d & e \\ a & 1, 0 & 1, 2 & 0, 1 \\ b & 0, 3 & 0, 1 & 2, 0 \end{array} \quad (1.12)$$

This game is solved in three rounds:

1. For Bob, d strictly dominates e , so delete e .
2. In the reduced game, a strictly dominates b for Alice, so delete b .
3. Against the remaining row a , Bob strictly prefers d to c , so delete c .

The sole surviving profile is (a, d) , with payoff $(1, 2)$. The successive deletions can be interpreted as increasingly strong assumptions about players' rationality and their reasoning about one another. A formal epistemic interpretation requires additional machinery.

Finite-game property

In a finite game, exhaustive elimination of strictly dominated pure strategies leaves the same surviving pure strategies regardless of the order of deletion. The corresponding statement is false for weak dominance: different deletion orders can leave different reduced games.

Many games, including matching pennies and coordination games, contain no dominated pure strategies. Dominance is therefore a useful first test rather than a universal solution method.

1.4 Nash equilibrium

Dominance asks whether a player can reject a strategy without predicting the opponents. Nash equilibrium instead asks whether all players' choices are mutually consistent. Equivalently, no player can gain through a unilateral deviation [2].

Definition

A profile $\mathbf{s}^* \in \mathbf{S}$ is a **pure-strategy Nash equilibrium** if

$$u_i(\mathbf{s}^*) \geq u_i(s_i, \mathbf{s}_{-i}^*) \quad \text{for every } i \in N \text{ and } s_i \in S_i. \quad (1.13)$$

Equilibrium is a stability condition. It does not imply uniqueness, fairness, or collective desirability, and it does not prevent players from benefiting through a coordinated change.

Nash equilibrium and dominance

If $\mathbf{s}^* = (s_1^*, \dots, s_n^*)$ is a strategy profile such that s_i^* is weakly dominant for every player i , then $\mathbf{s}^*$ is a Nash equilibrium.

Proof. Since s_i^* is weakly dominant,

$$u_i(s_i^*, \mathbf{s}_{-i}) \geq u_i(s_i, \mathbf{s}_{-i}) \quad (1.14)$$

for every $s_i \in S_i$ and every $\mathbf{s}_{-i} \in \mathbf{S}_{-i}$. In particular, for $\mathbf{s}_{-i} = \mathbf{s}_{-i}^*$,

$$u_i(s_i^*, \mathbf{s}_{-i}^*) \geq u_i(s_i, \mathbf{s}_{-i}^*) \quad (1.15)$$

for every player i and every $s_i \in S_i$. ◻

1.4.1 Best responses

Player i 's **best-response correspondence** is

$$\text{BR}_i(\mathbf{s}_{-i}) = \arg \max_{s_i \in S_i} u_i(s_i, \mathbf{s}_{-i}). \quad (1.16)$$

It is a set-valued mapping rather than necessarily a function because several strategies can tie for the maximum. In a finite game with nonempty strategy sets, the *best-response set* in ((1.16)) is nonempty.

Equivalent equilibrium test

$$\mathbf{s}^* \text{ is a Nash equilibrium} \Leftrightarrow s_i^* \in \text{BR}_i(\mathbf{s}_{-i}^*) \text{ for every } i. \quad (1.17)$$

Proof. This equivalence simply rewrites the inequalities in ((1.13)): each equilibrium strategy must maximize its player's utility against the opponents' equilibrium strategies. ◻

1.4.2 Finding pure equilibria

In each column, mark player 1's highest payoff. In each row, mark player 2's highest payoff. Include all ties. A cell is a pure-strategy Nash equilibrium exactly when both players' payoffs are marked.

Example: Bach or Stravinski

Two people want to spend the evening together, but they have different musical tastes: one prefers a Bach concert, the other prefers Stravinsky. Both would rather attend the same concert than go separately, but each prefers coordinating on their own favorite.

$$\begin{array}{cc}
 & \begin{array}{cc} B & S \end{array} \\
 \begin{array}{c} B \\ S \end{array} & \begin{array}{cc} 2, 1 & 0, 0 \\ 0, 0 & 1, 2 \end{array}
 \end{array} \tag{1.18}$$

It is easy to see that the only pure Nash equilibria are (B, B) and (S, S) .

Example: Coordination game

Two people want to meet, but they cannot communicate in advance. They must independently choose between two possible meeting places, say the station a or the café b . They mainly care about choosing the same place, but some location is preferred by both.

$$\begin{array}{cc}
 & \begin{array}{cc} a & b \end{array} \\
 \begin{array}{c} a \\ b \end{array} & \begin{array}{cc} 2, 2 & 0, 0 \\ 0, 0 & 1, 1 \end{array}
 \end{array} \tag{1.19}$$

The Nash equilibria are (a, a) and (b, b) .

The small games introduced above display several possible patterns:

Game	Nash equilibria	Main lesson
Prisoner's dilemma	(D, D)	a unique equilibrium can follow from dominance
Matching pennies	none	pure strategies may be insufficient
Sharing household chores	none	general-sum games may also lack a pure equilibrium
Bach or Stravinski	(B, B) and (S, S)	coordination with conflicting preferences
Coordination game	(a, a) and (b, b)	equilibrium need not select one outcome

No pure strategy is dominated in either game. In Bach or Stravinski, the players prefer different coordinated outcomes. In the coordination game, both prefer (a, a) to (b, b) , yet neither player can move from (b, b) to (a, a) unilaterally.

Matching pennies and the household-chores game show why the next chapter introduces mixed strategies. Allowing randomization restores a general existence result for finite games; this chapter makes no such claim for pure strategies.

1.5 Equilibrium and efficiency

The following concept tries to address the question *Which strategy profiles are socially efficient?* One possible answer is when no player's utility can be increased without simultaneously decreasing the utility to another player.

Definition

A profile $\mathbf{s}$ **Pareto dominates** a profile $\mathbf{t}$ if

$$u_i(\mathbf{s}) \geq u_i(\mathbf{t}) \quad \text{for every } i \in N, \quad (1.20)$$

with strict inequality for at least one player. A profile is **Pareto optimal** or **Pareto efficient** if no other profile Pareto dominates it.

Define the *social welfare* of strategy profile $\mathbf{s}$ by

$$w(\mathbf{s}) = \sum_{i \in N} u_i(\mathbf{s}). \quad (1.21)$$

Proposition

Every maximizer of w is Pareto optimal. Consequently, every finite normal-form game has at least one Pareto-optimal pure-strategy profile.

Proof. If $\mathbf{s}^*$ maximizes w and another profile $\mathbf{s}$ made every player weakly better off and at least one player strictly better off, summing the inequalities would give $w(\mathbf{s}) > w(\mathbf{s}^*)$, a contradiction. ◻

This is a method for finding one Pareto optimum, not a characterization of all Pareto optima. In a zero-sum game, every profile is Pareto optimal because the sum of utilities is constant.

The distinction from Nash equilibrium is fundamental:

Game	Strategic stability	Pareto efficiency
Prisoner's dilemma	(D, D) is the unique equilibrium	(D, D) is the only pure profile that is not Pareto optimal
Matching pennies	no pure equilibrium	every pure profile is Pareto optimal
Bach or Stravinski	(B, B) and (S, S) are equilibria	both equilibrium profiles are Pareto optimal
Coordination game	(a, a) and (b, b) are equilibria	only (a, a) is Pareto optimal

Nash equilibrium concerns unilateral deviations; Pareto optimality concerns joint comparisons. Neither concept generally implies the other.

1.6 Dynamic games and their normal form

The normal form can represent a dynamic interaction, but its strategies must then be understood as complete plans rather than single moves. The following Alice–Bob game makes the conversion explicit.

1.6.1 The Alice–Bob game

Alice and Bob are planning an evening together. Alice first decides whether they should stay home (a) or go out (b). If they stay home, Bob chooses between an action film (c) and cooking together (d). If they go out, Bob chooses a quiet cafe (e) or live music (f). In the latter case, Alice has the final choice between a rock club (g) and a jazz club (h). The terminal numbers are preference scores rather than monetary rewards.




Figure 1.1: Alice and Bob plan an evening. Payoffs are preference scores ordered (Alice, Bob).

Alice moves first. Bob observes her action and then moves at one of two decision nodes. If the history bf occurs, Alice moves again. Every player observes the complete preceding history, so the game has perfect information.

1.6.2 Strategies as complete contingent plans

A strategy prescribes an action at every decision point where the player might move, including decision points that will not be reached when that strategy is played. Therefore the strategy sets are

$$S_A = \{ag, ah, bg, bh\}, \quad S_B = \{ce, cf, de, df\}. \quad (1.22)$$

For example, ah instructs Alice to choose a initially and h if bf were reached. The second instruction is off the realized path when Alice chooses a , but it remains part of her strategy because a strategy must be a complete plan.

1.6.3 Constructing the normal form

To convert this game into the normal form, combine one complete plan for Alice with one complete plan for Bob, follow the path induced in Figure 1.1, and place the terminal payoff in the corresponding cell:

	<i>ce</i>	<i>cf</i>	<i>de</i>	<i>df</i>	
<i>ag</i>	0, 4	0, 4	2, 1	2, 1	
<i>ah</i>	0, 4	0, 4	2, 1	2, 1	(1.23)
<i>bg</i>	1, 2	3, 8	1, 2	3, 8	
<i>bh</i>	1, 2	8, 0	1, 2	8, 0	

Three cells illustrate the procedure:

- (ag, df) induces $a \rightarrow d$ and payoff $(2, 1)$;
- (bh, cf) induces $b \rightarrow f \rightarrow h$ and payoff $(8, 0)$;
- (bh, ce) induces $b \rightarrow e$ and payoff $(1, 2)$.

Repeated entries are not mistakes. Under ag , for instance, Alice chooses a immediately, so her instruction g can never be reached, while Bob's instruction after b is off the path for that profile. Both remain because a strategy is a complete contingent plan. Some off-path instructions affect play against other strategies; others, such as the difference between ag and ah , produce outcome-equivalent strategies and disappear only in a reduced normal form.

The normal-form game retains the utilities associated with every profile of complete plans, but it does not display which actions occur first, what is observed, or which instructions lie off the realized path. Different extensive-form games can consequently have the same normal form.

1.6.4 Nash equilibrium and sequential reasoning

Marking best responses in the matrix gives the unique pure-strategy Nash equilibrium

$$(bh, ce). \quad (1.24)$$

It induces the path $b \rightarrow e$ and payoff $(1, 2)$. The same outcome follows from solving the tree from its terminal decisions by *backward induction*:

1. after bf , Alice chooses h because $8 > 3$;
2. after b , Bob chooses e because $2 > 0$;
3. after a , Bob chooses c because $4 > 1$;
4. at the root, Alice chooses b because $1 > 0$.

Thus Alice proposes going out, but Bob chooses the cafe. Although both would prefer the rock-club outcome $(3, 8)$ to the cafe outcome $(1, 2)$, Alice would switch to her preferred jazz club once given the final choice. Bob anticipates that choice and avoids live music. The story illustrates how the inability to commit to a later action can rule out an outcome that both players prefer.

2. Mixed-strategy Nash equilibrium

2.1 Why randomize?

In [the previous chapter](#), matching pennies and the household-chores game had no pure-strategy Nash equilibrium. At every pure profile, someone could improve their payoff by changing their choice. This does not mean that these games have no equilibrium once players can randomize. It means that a fixed choice of a pure strategy may be too restrictive.

In matching pennies, a predictable choice can be exploited. A probability distribution over heads and tails allows a player to specify how they will randomize without making the realized choice predictable. The relevant question becomes: can the players choose distributions such that neither benefits from changing their own distribution?

This chapter develops the definitions and mathematical conditions needed to answer that question. Throughout, the normal-form game $(N, (S_i)_{i \in N}, (u_i)_{i \in N})$ is *finite*, that is, every S_i is a finite strategy set. We retain $S = \prod_{i \in N} S_i$ and the notation s_{-i} for the opponents' pure strategies.

2.2 Mixed strategies and expected utility

2.2.1 Probability distributions over pure strategies

Definition

A **mixed strategy** of player i is a probability distribution p_i on S_i . The set of mixed strategies is

$$\Delta_i = \left\{ p_i : S_i \rightarrow [0, 1] \left| \sum_{s_i \in S_i} p_i(s_i) = 1 \right. \right\}. \quad (2.1)$$

A **mixed-strategy profile** is $\mathbf{p} = (p_1, \dots, p_n) \in \Delta$, where $\Delta = \prod_{i \in N} \Delta_i$.

A player first selects a distribution and then draws a pure strategy according to it. The draws of different players are **independent**. Thus the probability of the pure profile $\mathbf{s}$ is

$$\prod_{i \in N} p_i(s_i). \quad (2.2)$$

Randomization does not require repeated play, although repeated independent draws make its frequencies easier to observe.

We identify a pure strategy s_i with the distribution assigning probability 1 to s_i and 0 otherwise. Consequently, pure strategies are special cases of mixed strategies. For example, if $S_i = \{H, T\}$, a mixed strategy can be written as $(p, 1 - p)$, with $0 \leq p \leq 1$. The endpoints are pure strategies. For $S_i = \{a, b, c\}$, the set Δ_i is a triangle, namely the **probability simplex** in Figure 2.1. Its vertices are pure strategies. Every edge contains distributions using at most two pure strategies. More generally, for $|S_i| = m$, the simplex has dimension $m - 1$.




Figure 2.1: A probability simplex for three pure strategies.

2.2.2 Expected utility

A rational player seeks to maximize expected utility.

Definition

The **expected utility** of player i at a mixed profile $\mathbf{p}$ is

$$U_i(\mathbf{p}) = \sum_{\mathbf{s} \in \mathbf{S}} u_i(\mathbf{s}) \prod_{j \in N} p_j(s_j). \quad (2.3)$$

We use the following notation: lowercase u_i evaluates a pure profile and uppercase U_i evaluates a mixed profile. In particular, against opponents' distributions $\mathbf{p}_{-i}$, the expected utility of a pure strategy s_i is

$$U_i(s_i, \mathbf{p}_{-i}) = \sum_{\mathbf{s}_{-i} \in \mathbf{S}_{-i}} u_i(s_i, \mathbf{s}_{-i}) \prod_{j \neq i} p_j(s_j). \quad (2.4)$$

Regrouping ((2.3)) by player i 's choice gives the identity used throughout this chapter:

$$U_i(p_i, \mathbf{p}_{-i}) = \sum_{s_i \in S_i} p_i(s_i) U_i(s_i, \mathbf{p}_{-i}). \quad (2.5)$$

For fixed opponents' strategies, a player's expected utility is a weighted average of the utilities of their pure strategies. In particular, it is linear in that player's own probabilities.

Note

Example 2.1. Player 1 chooses a row and player 2 chooses a column:

$$\begin{array}{cc|cc} & x & y & z \\ a & 1, 7 & 1, 5 & 3, 4 \\ b & 2, 3 & 0, 4 & 0, 6 \end{array} \quad (2.6)$$

Let $p = p_1(a)$, $q = p_2(x)$, and $r = p_2(y)$. Then $p_1(b) = 1 - p$ and $p_2(z) = 1 - q - r$. The expected utilities are

$$\begin{aligned} U_1(p_1, p_2) &= pq + pr + 3p(1 - q - r) + 2(1 - p)q, \\ U_2(p_1, p_2) &= 7pq + 5pr + 4p(1 - q - r) \\ &\quad + 3(1 - p)q + 4(1 - p)r + 6(1 - p)(1 - q - r). \end{aligned} \quad (2.7)$$

For player 1, the same calculation can be organized by pure strategies:

$$U_1(a, p_2) = q + r + 3(1 - q - r), \quad U_1(b, p_2) = 2q, \quad (2.8)$$

and $U_1(p_1, p_2) = pU_1(a, p_2) + (1 - p)U_1(b, p_2)$. We return to this game in Section 2.5.2 to see why an equilibrium may assign probability zero to some pure strategy.

2.3 Nash equilibrium and best responses

Definition

A mixed-strategy profile $\mathbf{p}^* \in \Delta$ is a **Nash equilibrium** if, for every player $i \in N$,

$$U_i(\mathbf{p}^*) \geq U_i(p_i, \mathbf{p}_{-i}^*) \quad \text{for every } p_i \in \Delta_i. \quad (2.9)$$

The comparison holds the opponents' distributions fixed. A deviation may change any of player i 's probabilities, including replacing randomization by a pure strategy. Equilibrium rules out a *strictly profitable* deviation.

2.3.1 The best-response correspondence

For fixed $\mathbf{p}_{-i}$, define

$$\text{BR}_i(\mathbf{p}_{-i}) = \operatorname{argmax}_{p_i \in \Delta_i} U_i(p_i, \mathbf{p}_{-i}). \quad (2.10)$$

The set-valued map BR_i is the **best-response correspondence**. Its value $\text{BR}_i(\mathbf{p}_{-i})$ is the set of mixed strategies of player i that are best responses to $\mathbf{p}_{-i}$. Directly from the definitions,

$$\mathbf{p}^* \text{ is a Nash equilibrium} \Leftrightarrow p_i^* \in \text{BR}_i(\mathbf{p}_{-i}^*) \text{ for every } i \in N. \quad (2.11)$$

2.3.2 Why pure strategies suffice to test Nash equilibrium

Although there are infinitely many mixed strategies, there are only finitely many pure strategies to test if a given mixed strategy profile is an equilibrium.

Proposition

A profile $\mathbf{p}^*$ is a Nash equilibrium if and only if, for every player i and every $s_i \in S_i$,

$$U_i(\mathbf{p}^*) \geq U_i(s_i, \mathbf{p}_{-i}^*). \quad (2.12)$$

Proof. Necessity follows because pure strategies are mixed strategies. Conversely, suppose all the displayed inequalities hold. For any mixed strategy p_i , identity ((2.5)) gives

$$\begin{aligned}
U_i(p_i, \mathbf{p}_{-i}^*) &= \sum_{s_i \in S_i} p_i(s_i) U_i(s_i, \mathbf{p}_{-i}^*) \\
&\leq \sum_{s_i \in S_i} p_i(s_i) U_i(\mathbf{p}^*) = U_i(\mathbf{p}^*) \sum_{s_i \in S_i} p_i(s_i) = U_i(\mathbf{p}^*).
\end{aligned} \tag{2.13}$$

□

2.4 Existence of equilibrium

Nash's theorem [2]

Every finite normal-form game has at least one equilibrium in mixed strategies.

The proof uses a fixed-point theorem. We use the existence result here without reproducing that proof. Several qualifications matter. A mixed strategy includes a pure strategy, so the theorem does not assert that some player must randomize. It provides only an existence guarantee. The methods for finding an equilibrium belong to the later topic on computing Nash equilibrium.

2.5 Supports and indifference

Definition

The **support** of a mixed strategy p_i is the set of pure strategies

$$\text{spt}(p_i) = \{s_i \in S_i \mid p_i(s_i) > 0\}. \tag{2.14}$$

The mixed strategy is **completely mixed** if its support is S_i . A profile $(p_1, \dots, p_n)$ is completely mixed if every p_i is.

For example, let $S_i = \{a, b, c\}$ and $p_i(a) = p_i(b) = p_i(c) = \frac{1}{3}$, then p_i is completely mixed. Another mixed strategy $q_i(a) = q_i(b) = \frac{1}{2}$ has support $\text{spt}(q_i) = \{a, b\}$. A pure strategy has a one-element support.

2.5.1 The indifference principle

Proposition

At a Nash equilibrium $\mathbf{p}^*$, every pure strategy $s_i \in \text{spt}(p_i^*)$ gives player i the same expected utility:

$$U_i(s_i, \mathbf{p}_{-i}^*) = U_i(\mathbf{p}^*). \tag{2.15}$$

In particular, any two pure strategies in the support have equal expected utilities against the equilibrium opponents' strategies.

Proof. Suppose $s_i, t_i \in \text{spt}(p_i^*)$ and $U_i(s_i, \mathbf{p}_{-i}^*) > U_i(t_i, \mathbf{p}_{-i}^*)$. Transfer all probability assigned to t_i to s_i , leaving the other probabilities unchanged. Specifically, define:

$$p_i(a) = \begin{cases} 0, & a = t_i, \\ p_i^*(s_i) + p_i^*(t_i), & a = s_i, \\ p_i^*(a), & a \notin \{s_i, t_i\}. \end{cases} \tag{2.16}$$

This is a mixed strategy. Its improvement in expected utility is

$$U_i(p_i, \mathbf{p}_{-i}^*) - U_i(\mathbf{p}^*) = p_i^*(t_i) [U_i(s_i, \mathbf{p}_{-i}^*) - U_i(t_i, \mathbf{p}_{-i}^*)] > 0, \tag{2.17}$$

contradicting equilibrium. Therefore the utilities on the support are equal. Their weighted average is $U_i(\mathbf{p}^*)$, so their common value must be that expected utility. □

The reason for mixing is therefore not that a weighted average improves on the best pure reply. *A player mixes among equally good replies, and the chosen probabilities in turn affect the opponents' incentives.*

Example: Indifference is not sufficient

Player 1 chooses a row and player 2 chooses a column in the following 3×2 game:

$$\begin{array}{cc} & \begin{array}{c} l \\ r \end{array} \\ \begin{array}{c} a \\ b \\ c \end{array} & \begin{array}{cc} 1, 0 & 0, 1 \\ 0, 1 & 1, 0 \\ 0.6, 0 & 0.6, 0 \end{array} \end{array} \quad (2.18)$$

Suppose player 1 uses only a, b , with probabilities $p, 1 - p$, and player 2 uses l, r , with probabilities $q, 1 - q$. Indifference within these proposed supports requires

$$\begin{aligned} U_1(a, p_2) &= q = 1 - q = U_1(b, p_2), \\ U_2(p_1, l) &= 1 - p = p = U_2(p_1, r). \end{aligned} \quad (2.19)$$

Thus $p = q = 1/2$, giving the candidate profile

$$p_1 = (1/2, 1/2, 0), \quad p_2 = (1/2, 1/2), \quad (2.20)$$

with coordinates ordered as (a, b, c) and (l, r) . Every strategy in each player's support gives that player expected utility $1/2$. Nevertheless, player 1 can switch to the unused strategy c and obtain

$$U_1(c, p_2) = 0.6 > \frac{1}{2} = U_1(p_1, p_2). \quad (2.21)$$

The candidate mixed strategy is therefore *not a Nash equilibrium*.

In conclusion, indifference among strategies in the support is *necessary but not sufficient* for equilibrium: we must also check that no strategy outside the support gives a higher payoff.

2.5.2 The support characterization

Indifference among used pure strategies is necessary for an equilibrium. The question remains if an unused strategy outside the support might still be better. To state the necessary and sufficient condition for equilibrium, distinguish pure from mixed best responses:

$$B_i(\mathbf{p}_{-i}) = \operatorname{argmax}_{s_i \in S_i} U_i(s_i, \mathbf{p}_{-i}). \quad (2.22)$$

The set $B_i(\mathbf{p}_{-i})$ consists of pure strategies, whereas $\operatorname{BR}_i(\mathbf{p}_{-i})$ consists of probability distributions. Since S_i is finite, $B_i(\mathbf{p}_{-i}) \neq \emptyset$.

Proposition

A profile of mixed strategies $\mathbf{p}^*$ is a Nash equilibrium if and only if

$$\operatorname{spt}(p_i^*) \subseteq B_i(\mathbf{p}_{-i}^*) \quad \text{for every } i \in N. \quad (2.23)$$

Proof. Let $\mathbf{p}^*$ be an equilibrium. For any pure strategy $s_i \in \operatorname{spt}(p_i^*)$, the indifference principle says that

$$U_i(s_i, \mathbf{p}_{-i}^*) = U_i(\mathbf{p}^*). \quad (2.24)$$

For any pure strategy $s_i \in S_i \setminus \operatorname{spt}(p_i^*)$, the pure deviation condition is

$$U_i(s_i, \mathbf{p}_{-i}^*) \leq U_i(\mathbf{p}^*) \quad (2.25)$$

Therefore,

$$U_i(\mathbf{p}^*) = \max_{s_i \in S_i} U_i(s_i, \mathbf{p}_{-i}^*), \quad (2.26)$$

which proves inclusion ((2.23)).

Conversely, let ((2.23)) hold for a profile $\mathbf{p}^*$. Then

$$\begin{aligned} U_i(\mathbf{p}^*) &= \sum_{s_i \in \text{spt}(p_i^*)} p_i^*(s_i) U_i(s_i, \mathbf{p}_{-i}^*) = \sum_{s_i \in \text{spt}(p_i^*)} p_i^*(s_i) \max_{t_i \in S_i} U_i(t_i, \mathbf{p}_{-i}^*) \\ &= \max_{t_i \in S_i} U_i(t_i, \mathbf{p}_{-i}^*) \cdot \sum_{s_i \in \text{spt}(p_i^*)} p_i^*(s_i) = \max_{t_i \in S_i} U_i(t_i, \mathbf{p}_{-i}^*), \end{aligned} \quad (2.27)$$

where ((2.23)) was used in the second equality. The derived identity

$$U_i(\mathbf{p}^*) = \max_{t_i \in S_i} U_i(t_i, \mathbf{p}_{-i}^*) \quad (2.28)$$

means that $\mathbf{p}^*$ is a Nash equilibrium. ◻

Equivalently, condition ((2.23)) says that for every player i :

- every strategy in the support gives expected utility $U_i(\mathbf{p}^*)$;
- every strategy outside the support gives expected utility at most $U_i(\mathbf{p}^*)$.

Example: An equilibrium with partial support

Return to the 2×3 game in [Example 2.1](#). Consider mixed strategies

$$p_1^* = (1/2, 1/2), \quad p_2^* = (3/4, 0, 1/4), \quad (2.29)$$

with coordinates ordered as (a, b) and (x, y, z) . Against p_2^* , player 1 obtains

$$U_1(a, p_2^*) = \frac{3}{4} + 3 \cdot \frac{1}{4} = \frac{3}{2}, \quad U_1(b, p_2^*) = 2 \cdot \frac{3}{4} = \frac{3}{2}. \quad (2.30)$$

Against p_1^* , player 2 obtains

$$U_2(p_1^*, x) = 5, \quad U_2(p_1^*, y) = \frac{9}{2}, \quad U_2(p_1^*, z) = 5. \quad (2.31)$$

Both players mix only among best responses, so this profile is a Nash equilibrium with expected payoffs $(3/2, 5)$. Player 2 leaves y unused because its expected payoff is strictly lower.

The inclusion ((2.23)) need not be equality: an unused pure strategy outside the support is allowed to tie with those in the support. If the profile is completely mixed, there are no unused strategies, so indifference for every player is sufficient.

2.6 Two-player (2×2) examples

This case is elementary. Let p and q denote the players' probabilities of choosing their first strategies. The indifference principle gives two linear equations: solve for q by making the row player indifferent, and for p by making the column player indifferent. If $0 < p, q < 1$, they define a completely mixed strategy profile and hence an equilibrium.

Example: Matching pennies

Alice chooses the row and wants the coins to match; Bob chooses the column and wants them to differ:

$$\begin{array}{cc} & \begin{array}{cc} H & T \end{array} \\ \begin{array}{c} H \\ T \end{array} & \begin{array}{cc} 1, -1 & -1, 1 \\ -1, 1 & 1, -1 \end{array} \end{array} \quad (2.32)$$

Let $p = \Pr(\text{Alice chooses } H)$ and $q = \Pr(\text{Bob chooses } H)$. Alice's pure-strategy utilities are

$$U_A(H, q) = 2q - 1, \quad U_A(T, q) = 1 - 2q. \quad (2.33)$$

She prefers H when $q > 1/2$, prefers T when $q < 1/2$, and is indifferent when $q = 1/2$. Bob's utilities are

$$U_B(p, H) = 1 - 2p, \quad U_B(p, T) = 2p - 1, \quad (2.34)$$

so he is indifferent when $p = 1/2$. At

$$p^* = q^* = \frac{1}{2}, \quad (2.35)$$

every pure strategy gives its player expected utility 0. It is the unique equilibrium: if either player chose a non-uniform distribution, the other would have a unique pure best response, which would in turn invite a profitable change by the first player.

Example: Sharing household chores

Alice chooses whether to cook or clean, and Bob chooses whether to shop or take out the trash. Alice prefers chores that contribute to the same goal; Bob prefers covering different needs. Their satisfaction scores are:

$$\begin{array}{cc} & \begin{array}{cc} \text{Shop} & \text{Trash} \end{array} \\ \begin{array}{c} \text{Cook} \\ \text{Clean} \end{array} & \begin{array}{cc} 2, 0 & 0, 1 \\ 0, 2 & 1, 0 \end{array} \end{array} \quad (2.36)$$

Let $p = \Pr(\text{Cook})$ and $q = \Pr(\text{Shop})$. Alice's indifference requires

$$2q = 1 - q, \quad \text{hence } q = \frac{1}{3}. \quad (2.37)$$

Bob's indifference requires

$$2(1 - p) = p, \quad \text{hence } p = \frac{2}{3}. \quad (2.38)$$

This mixed strategy profile is a unique equilibrium with expected payoffs $(2/3, 2/3)$. Notice the direction of the equations: Alice's indifference determines how often *Bob* shops, and Bob's indifference determines how often *Alice* cooks.

Example: Bach or Stravinski

Alice prefers Bach and Bob prefers Stravinski, but both prefer attending together to attending different concerts:

$$\begin{array}{cc}
 & B & S \\
 B & 2, 1 & 0, 0 \\
 S & 0, 0 & 1, 2
 \end{array} \tag{2.39}$$

The pure equilibria are (B, B) and (S, S) . Let $p = \Pr(\text{Alice chooses } B)$ and $q = \Pr(\text{Bob chooses } B)$. For both to mix, Alice must be indifferent between $2q$ and $1 - q$, while Bob must be indifferent between p and $2(1 - p)$. Therefore there is also a mixed equilibrium with

$$p = \frac{2}{3}, \quad q = \frac{1}{3}, \tag{2.40}$$

and the expected utilities are $(\frac{2}{3}, \frac{2}{3})$. Both pure equilibria give both players more than the mixed equilibrium. Nevertheless, neither player can improve unilaterally at the mixed profile.

Example: Coordination

Players independently choose a meeting place. Both prefer meeting at a to meeting at b , and both prefer either meeting to missing one another:

$$\begin{array}{cc}
 & a & b \\
 a & 2, 2 & 0, 0 \\
 b & 0, 0 & 1, 1
 \end{array} \tag{2.41}$$

Besides (a, a) and (b, b) , there is an equilibrium in which both choose a with probability $1/3$. Indeed, against that distribution, choosing a gives $2/3$ and choosing b gives $1 - 1/3 = 2/3$. The expected payoff pair is $(2/3, 2/3)$.

Example: Professor's Dilemma

This example is adapted from the *Professor's Dilemma* in Procaccia's lecture slides [10, slide 8]. A professor decides whether to prepare a lecture, while the students, modelled as one player, decide whether to listen. They choose without observing the other's decision. A prepared lecture with attentive students benefits both. Effort without participation costs the player who makes it.

$$\begin{array}{cc}
 & \text{Listen} & \text{Sleep} \\
 \text{Prepare} & 10^6, 10^6 & -10, 0 \\
 \text{Slack off} & 0, -10 & 0, 0
 \end{array} \tag{2.42}$$

Both (Prepare, Listen) and (Slack off, Sleep) are pure equilibria. If the students listen with probability q , preparing gives $10^6q - 10(1 - q)$, while slacking off gives 0. The professor is indifferent when

$$q = \frac{10}{10^6 + 10} = \frac{1}{100001}. \tag{2.43}$$

The students' analogous indifference condition gives the same probability for preparation. Hence a completely mixed equilibrium also exists, with expected payoff 0 to each player. The very large benefit from successful coordination does not remove the other equilibria or explain which one will be played.

2.7 Good Samaritan: Who will provide aid?

This example is adapted from the *Good Samaritan Game* in Chapter 2 of Binmore's *Game Theory: A Very Short Introduction* [11, pp. 25–26]. There are $n \geq 2$ bystanders who hear a cry for aid. Each independently chooses whether to provide aid (a) or leave the task to someone else ($\bar{a}$), without observing the others' choices. Everyone values someone receiving help at 10 utility units. Providing aid costs 1 utility unit, so a helper receives 9 regardless of how many others also help. A non-helper receives 10 if someone else helps and 0 if nobody helps.

Thus $S_i = \{a, \bar{a}\}$ and

$$u_i(\mathbf{s}) = \begin{cases} 9, & s_i = a, \\ 10, & s_i = \bar{a} \text{ and some } j \neq i \text{ chooses } a, \\ 0, & \mathbf{s} = (\bar{a}, \dots, \bar{a}). \end{cases} \quad (2.44)$$

For $n = 2$, the game is

$$\begin{array}{cc} & \begin{matrix} a & \bar{a} \end{matrix} \\ \begin{matrix} a \\ \bar{a} \end{matrix} & \begin{pmatrix} 9, 9 & 9, 10 \\ 10, 9 & 0, 0 \end{pmatrix} \end{array} \quad (2.45)$$

2.7.1 Pure equilibria

There are exactly n pure equilibria, each with *exactly one helper*.

- If nobody helps, any bystander can increase their utility from 0 to 9 by helping.
- If at least two bystanders help, any helper can stop helping and increase their utility from 9 to 10.
- If exactly one helps, that helper would lose 9 by stopping, while every other bystander would lose 1 by also helping.

Pure equilibria therefore exist, but they assign different roles to identical players. The model by itself does not select the helper.

2.7.2 A symmetric mixed equilibrium

We now look for a particular equilibrium in which everyone uses the same distribution. This is a restriction on the equilibrium we seek, not a consequence that every equilibrium must satisfy. Let q be each bystander's probability of *not* helping and $\mathbf{p}$ be the corresponding strategy profile. Then

$$U_i(a, \mathbf{p}_{-i}) = 9, \quad U_i(\bar{a}, \mathbf{p}_{-i}) = 10 (1 - q^{n-1}). \quad (2.46)$$

For both pure strategies to be used with positive probability, indifference requires

$$9 = 10 (1 - q^{n-1}), \quad q = (0.1)^{1/(n-1)}. \quad (2.47)$$

Since $0 < q < 1$, this gives a completely mixed profile. This shows that the profile is a Nash equilibrium with

$$U_i(\mathbf{p}^*) = 9 \quad \text{for every } i \in N. \quad (2.48)$$

2.7.3 The chance that nobody helps

There are $n - 1$ other bystanders in an individual's indifference condition, but there are n bystanders who must all abstain for aid to fail. Consequently, at the symmetric equilibrium,

$$\Pr(\text{nobody helps}) = q^n = (0.1)^{n/(n-1)}. \quad (2.49)$$

Table 2.1: Symmetric equilibrium probabilities, rounded to three decimal places.

Bystanders n	Each player's probability of helping $1 - q$	Probability nobody helps q^n
2	0.900	0.010
3	0.684	0.032
5	0.438	0.056
10	0.226	0.077

As $n \rightarrow \infty$,

$$q \rightarrow 1, \quad 1 - q \rightarrow 0, \quad q^n \rightarrow 0.1. \quad (2.50)$$

Each individual becomes less likely to help as the group grows. The chance of no aid increases toward 10% in this family of symmetric equilibria, even though there are more potential helpers. This is a consequence of the specified incentives and independent equilibrium randomization, rather than an empirical claim about actual emergencies.

Appendices

Normal-form game catalogue

All reusable normal-form games are stored together in `_variables.yml`. A lecture inserts a selected matrix with a Quarto variable shortcode, for example `games.prisoners_dilemma.matrix`. The catalogue below renders those same definitions.

Instigation game

$$\begin{array}{cc} & s & \neg s \\ a & d, -t & -v, m \\ \neg a & 0, r & 0, 0 \end{array} \quad (1)$$

Prisoner's dilemma

$$\begin{array}{cc} & C & D \\ C & -1, -1 & -4, 0 \\ D & 0, -4 & -3, -3 \end{array} \quad (2)$$

Matching pennies

$$\begin{array}{cc} & H & T \\ H & 1, -1 & -1, 1 \\ T & -1, 1 & 1, -1 \end{array} \quad (3)$$

Sharing household chores

$$\begin{array}{cc} & \text{Shop} & \text{Trash} \\ \text{Cook} & 2, 0 & 0, 1 \\ \text{Clean} & 0, 2 & 1, 0 \end{array} \quad (4)$$

Bach or Stravinski

$$\begin{array}{cc} & B & S \\ B & 2, 1 & 0, 0 \\ S & 0, 0 & 1, 2 \end{array} \quad (5)$$

Coordination game

$$\begin{array}{cc} & a & b \\ a & 2, 2 & 0, 0 \\ b & 0, 0 & 1, 1 \end{array} \quad (6)$$

Iterated-dominance example

$$\begin{array}{ccc} & c & d & e \\ a & 1, 0 & 1, 2 & 0, 1 \\ b & 0, 3 & 0, 1 & 2, 0 \end{array} \quad (7)$$

Strategic form of the dynamic game

	<i>ce</i>	<i>cf</i>	<i>de</i>	<i>df</i>	
<i>ag</i>	0, 4	0, 4	2, 1	2, 1	(8)
<i>ah</i>	0, 4	0, 4	2, 1	2, 1	
<i>bg</i>	1, 2	3, 8	1, 2	3, 8	
<i>bh</i>	1, 2	8, 0	1, 2	8, 0	

Expected utility and partial support

	<i>x</i>	<i>y</i>	<i>z</i>	
<i>a</i>	1, 7	1, 5	3, 4	(9)
<i>b</i>	2, 3	0, 4	0, 6	

Indifference is not sufficient

	<i>l</i>	<i>r</i>	
<i>a</i>	1, 0	0, 1	(10)
<i>b</i>	0, 1	1, 0	
<i>c</i>	0.6, 0	0.6, 0	

Penalty kicks

Rows are the kicker's choices and columns are the goalkeeper's choices. Entries give scoring probabilities; the goalkeeper's utility is their negative.

	<i>u</i>	<i>n</i>	
<i>u</i>	0.58	0.95	(11)
<i>n</i>	0.93	0.70	

Professor's dilemma

	Listen	Sleep	
Prepare	$10^6, 10^6$	$-10, 0$	(12)
Slack off	$0, -10$	$0, 0$	

Good Samaritan: two bystanders

Action *a* means providing aid and $\bar{a}$ means not providing aid.

	<i>a</i>	$\bar{a}$	
<i>a</i>	9, 9	9, 10	(13)
$\bar{a}$	10, 9	0, 0	

Extensive-form game catalogue

Alice–Bob game

The game tree from [Strategic games](#).




Figure -1: Alice and Bob plan an evening. Payoffs are preference scores ordered (Alice, Bob).

References

- [1] J. von Neumann, “Zur Theorie der Gesellschaftsspiele,” *Mathematische Annalen*, vol. 100, pp. 295–320, 1928, doi: [10.1007/BF01448847](https://doi.org/10.1007/BF01448847).
- [2] J. Nash, “Non-Cooperative Games,” *Annals of Mathematics*, vol. 54, no. 2, pp. 286–295, 1951, doi: [10.2307/1969529](https://doi.org/10.2307/1969529).
- [3] L. S. Shapley, “A Value for n-Person Games,” *Contributions to the Theory of Games II*. Princeton University Press, Princeton, pp. 307–317, 1953.
- [4] W. Vickrey, “Counterspeculation, Auctions, and Competitive Sealed Tenders,” *The Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961, doi: [10.1111/j.1540-6261.1961.tb02789.x](https://doi.org/10.1111/j.1540-6261.1961.tb02789.x).
- [5] R. B. Myerson, “Optimal Auction Design,” *Mathematics of Operations Research*, vol. 6, no. 1, pp. 58–73, 1981, doi: [10.1287/moor.6.1.58](https://doi.org/10.1287/moor.6.1.58).
- [6] D. Silver *et al.*, “Mastering the Game of Go with Deep Neural Networks and Tree Search,” *Nature*, vol. 529, pp. 484–489, 2016, doi: [10.1038/nature16961](https://doi.org/10.1038/nature16961).
- [7] M. Moravčík *et al.*, “DeepStack: Expert-Level Artificial Intelligence in Heads-Up No-Limit Poker,” *Science*, vol. 356, no. 6337, pp. 508–513, 2017, doi: [10.1126/science.aam6960](https://doi.org/10.1126/science.aam6960).
- [8] O. Vinyals *et al.*, “Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement Learning,” *Nature*, vol. 575, pp. 350–354, 2019, doi: [10.1038/s41586-019-1724-z](https://doi.org/10.1038/s41586-019-1724-z).
- [9] V. Kovařík, E. O. Chen, S. Petersen, A. Ghersengorin, and V. Conitzer, “AI Testing Should Account for Sophisticated Strategic Behaviour.” 2025. doi: [10.48550/arXiv.2508.14927](https://doi.org/10.48550/arXiv.2508.14927).
- [10] A. D. Procaccia, “Noncooperative Games 1: Basic Concepts.” Accessed: Sept. 27, 2026. [Online]. Available: <https://www.cs.cmu.edu/~arielpro/15896s16/slides/896s16-17.pdf>
- [11] K. Binmore, *Game Theory: A Very Short Introduction*. Oxford University Press, 2007. doi: [10.1093/actrade/9780199218462.001.0001](https://doi.org/10.1093/actrade/9780199218462.001.0001).