

Chapter 1 of “Potential Game Theory”

Theories and applications

Author: Robert P. Gilles

Institute: Department of Economics

The Queen’s University of Belfast

Date: April 17, 2026

Version: 2.00

Email: r.gilles@qub.ac.uk

To the many shoulders of thinkers on which this work rests

Contents

Chapter 1 Games: The normal and potential forms	2
1.1 What is game theory?	2
1.2 Formalising normal form games	5
1.2.1 Pre-games: Players and actions	5
1.2.2 Normal form games: Payoff structures	10
1.2.3 Nash equilibrium	14
1.3 Introducing potentials: Potential form games	22
1.3.1 Defining equilibrium in potential form games	23
1.3.2 Nash potential games	27
1.4 Notes	31
1.4.1 A short history of potentials in game theory	32
1.4.2 Main fields of application	33
1.5 Proofs of Theorems and Propositions	34
1.5.1 Proof of Theorem 1.2	34
1.5.2 Proof of Theorem 1.3	35
1.5.3 Proof of Theorem 1.4	36
1.5.4 Proof of Theorem 1.5	38
Bibliography	40

Meaning of the word “potential”.

Meaning 1: Existing in possibility, not in act.

Meaning 2: Having the effect without the external actual property.

Meaning 3: Efficacious; powerful. (Noted as “not in use”.)

Meaning 4: Grammatical usage – a mood denoting possibility.

Description taken from Samuel Johnson’s *A Dictionary of the English Language* (1755).

Preface and acknowledgements

Preface yet to be written.

Chapter 1

Games: The normal and potential forms

Chapter summary

- Definition of non-cooperative games*
- Normal form games*
- Nash equilibrium*
- Existence of Nash equilibrium*
- Potential functions*
- Potential form games*
- Equilibrium in potential form games*
- Characterisation of equilibria*
- Definition of Nash potential games*
- Equivalence theorem for Nash potentials*

1.1 What is game theory?

Game theory can be understood as a set of tools that enables the description and analysis of social interactive decision situations. This perspective has several important implications. Firstly, game theory is not a single unified theory or research programme in the sense of Lakatos (1978), but rather a collection of various subfields and approaches. Each subfield represents a distinct approach to describing specific social interactive decision situations.

The common subject matter of social interactive decision making unifies these diverse subfields and yields a shared understanding of what constitutes an interactive decision situation. The following common features characterise such situations:

- (i) Multiple decision makers are involved within a social decision situation. Each decision maker has the potential to control specific decisions within that situation.¹
- (ii) Each decision maker possesses a set of actions from which they can choose when making a decision. These potential choices are commonly referred to as actions. Thus, at each decision moment, a decision maker has control over a range of multiple actions to select from.²
- (iii) The decisions made by different decision makers interact with one another, influencing the resulting outcome of the social decision situation. If any decision maker alters their selected action, it potentially affects the outcome.

Based on these shared features, we can introduce established game-theoretic terminology. A social interactive decision situation is known as a *game*, with the decision makers referred to as *players*. Each player is assumed to have control over at least one decision moment within the course of the game. At these decision moments, players select *actions* from a well-defined set of multiple options.

I explore next some well-established mathematical game-theoretic frameworks in which these interactive decision situations are represented.

¹It is clear that if we have only one decision maker (player), then there is simply no interaction possible. Instead one then arrives at a standard decision problem, which the subject of study in mathematical decision theory.

²If a decision maker would only control a set consisting of a single action for a certain decision moment, she would actually not make any decision. She could therefore be omitted as a relevant decision maker at that moment in the decision situation.

The fundamental approaches in game theory There are two fundamentally different approaches to the description of an interactive decision situation. The first approach is based on the absolute absence of any binding agreements between the decision makers in these interactive decision situations. This is also known as *non-cooperative game theory*.

The non-cooperative approach fits very well with many applications of a social interactive decision situation as described by a game. Indeed, it assumes that each player in a game is driven by a well-formulated goal. This goal is formalised as the player's payoff function. This function assigns to each outcome resulting from a selection of individual actions a payoff. Each player now optimises her payoff by selecting actions that are under her control. How this is determined is actually the subject matter of non-cooperative game theory. As such, non-cooperative game theory is the most pristine expression of the principle of *methodological individualism* that lies at the foundation of most of contemporary microeconomics.

The second fundamental approach is known as cooperative game theory and allows players to write binding contracts. This changes the analysis and interpretation of a game radically. Indeed, if binding agreements can be written, all players collectively will pursue the maximisation of the total wealth that can be generated within the social decision situation at hand. A binding contract then determines how this generated wealth is distributed among the various players in this interactive decision situation.

Thus, the main objective of cooperative game theory is to determine a “just” or “well-supported” contract between all players to divide the total wealth generated collectively. Such a contract can be based on pure bargaining power or solely on fairness considerations or mixtures of both power and fairness.

Although non-cooperative and cooperative game theory are fundamentally separated through the acceptance of the hypothesis of allowing binding contracts, in technical terms the various representations of interactive decision situations span these two approaches.

Normal form non-cooperative games There are two main conceptions to represent non-cooperative interactive decision making: the normal form and the extensive form.

The fundamental method of description that informs our analysis in this text is that of the *normal form game*. A normal form game is a mathematical representation of direct strategic interaction among multiple players, where each player makes decisions independently without any explicit communication or coordination. Each player has a set of actions—also referred to as “strategies” in much of the literature—from which she chooses a complete plan specifying what to do in every contingency. Associated with each profile of chosen actions is a payoff for every player, which may be expressed in monetary values, points, or any other relevant measure.

Decisions are made simultaneously, without knowledge of the choices made by the other players. Players are assumed to have complete information about the available actions and payoffs, but may have imperfect or asymmetric information about each others' preferences and past actions. The concept of Nash equilibrium is central to this framework: a Nash equilibrium is a profile at which no player can unilaterally change her action to improve her own payoff, given the actions chosen by the others. It represents a stable state of the game at which no player has an incentive to deviate.

The normal form representation, despite its limitations in comparison with the extensive form representation, has a significant advantage in that the analysis is much more transparent. In particular, few equilibrium concepts beyond Nash equilibrium are meaningful in this setting, so that more powerful insights can be established within the normal form.

Extensive form games The normal form representation stands in contrast to the second conception, the *extensive form game*. An extensive form game represents the sequential nature of decision-making: players take turns to make decisions, and the timing and order of their actions are explicitly represented by a tree-like structure called a *game tree*. A game tree is a graphical representation consisting of nodes—decision points—connected by branches that represent the possible actions available at each decision point. The tree starts with a single node called the root, representing the initial state of the game, and branches out to subsequent nodes as players make decisions.

Players may have imperfect or asymmetric information about the actions taken by previous players. To capture this, *information sets* are used to group together nodes that are indistinguishable to a player based on their available information at that point, representing a player's lack of knowledge about which node they are at within the information set. In an extensive form game, players choose *strategies* rather than specific actions at each decision point; a strategy represents a complete plan of action, specifying what choice the player will make at every decision point, given their available information. As in a normal form game, there are payoffs or utilities associated with each possible outcome.

An extensive form game can be further classified into *perfect information* and *imperfect information* games. In a perfect information game, each player knows the exact history of actions and can observe the actions taken by previous players. In an imperfect information game, players may have private information or uncertain knowledge about the actions taken by previous players.

These defining features of an extensive form game allow for the modelling of sequential decision-making, the representation of imperfect and asymmetric information, and the analysis of strategies that involve both actions and timing. Game theory uses these features to study strategic interaction and predict likely outcomes in such games.

Several equilibrium concepts extend and complement Nash equilibrium in the extensive form. *Subgame perfect equilibrium* (SPE) requires that strategies constitute a Nash equilibrium in every subgame and thereby captures optimality at every decision point, not only the initial ones. *Sequential equilibrium* refines SPE by requiring optimality off the equilibrium path as well, under a consistency condition on players' beliefs. *Perfect Bayesian equilibrium* (PBE) combines sequential rationality with Bayesian updating, requiring strategies to be sequentially rational at every information set and consistent with players' beliefs about the actions of the others and the moves of nature. These refinements provide more discriminating predictions than Nash equilibrium by accounting for the sequential nature of decision-making, information sets, and players' beliefs.

Potential form games In this text we focus on normal form games; within this class we identify specific subclasses known as potential form games, in which the payoff structure can be summarised through a single scalar function—the *potential function*—defined on the profile space.

There are several distinct ways in which the payoff structure can be summarised through such a function, giving rise to different sub-classes of potential form games. In an *exact potential game*, the potential function reproduces each player's payoff differences under unilateral deviations exactly. In an *ordinal potential game*, the potential function only preserves the ordering of payoff differences: a unilateral deviation that increases a player's payoff also increases the potential, but the magnitudes need not coincide. A player in an exact potential game can therefore read off the exact magnitude of her improvement from the potential, whereas a player in an ordinal potential game can read off only its sign.

Sources on the theory of games and game forms There are numerous textbooks on game theory and associated solution concepts. The field originated in the seminal work of Neumann and Morgenstern (1947), which defined the field and the conceptions that still determine this field of scientific pursuit.

For textbook treatments, I restrict myself to the main established texts. A good textbook for those unfamiliar with the subject I refer to Osborne (2004). For more advanced treatments of the subject I refer to Fudenberg and Tirole (1991), Myerson (1991), Osborne and Rubinstein (1994), Owen (2013), Vega-Redondo (2003), and Maschler, Solan, and Zamir (2013). There is no dedicated textbook treatment of the subject of potential form games.

1.2 Formalising normal form games

In our exploration of potential games we confine our discussion to the normal form representation. The normal form is the most straightforward mathematical model of an interactive decision situation involving rational agents. It combines a “pre-game”—which enumerates the decision makers (the “players”) and their available actions—with a payoff structure that assigns to every profile of actions the resulting utilities, payouts, or rewards to the players. A normal form game therefore consists of three elements: (i) the players; (ii) the players’ action sets; and (iii) the payoff structure.

We restrict ourselves mainly to “finite” games, where decision situations are described that involve a given finite number of decision makers. In specifically denoted sections, we might consider continuum games, which refer to theoretical constructs in which a large mass of decision makers are involved. These sections will be indicated specifically and explicitly throughout this text.

1.2.1 Pre-games: Players and actions

Game theory is firmly rooted in the hypothesis that in interactive decision situations, the described decision makers are *intelligent* or *rational*. This means that decisions are made considerately and deliberately. In fact, throughout we assume that a normal form game is a model of interactive decision-making in which each decision-maker chooses his plan of action once and for all, and these choices are made simultaneously; that these choices are guided by aims or the achievement of goals; and that decision makers are acting prescriptively and follow a specified model of decision making.

An intelligent or rational decision maker is referred to as a *player*. A player is essentially a mathematical point or node to which can be assigned elements that describe the decision making processes related to that decision maker. Formally, players are introduced through a **player set**:

$$N = \{1, \dots, n\} \text{ with } n \in \mathbb{N}. \quad (1.1)$$

Here, $n \in \mathbb{N}$ denotes the number of players or decision makers involved in the decision situation under consideration. The player set is a defining, specific setting on which the decision situation is described.

The next element of the model is which decisions a particular decision maker can select from. It is assumed that each player is assigned a specified set of actions from which she can choose. Throughout we assume that each player has at least *two* potential actions to choose from. Also, it is assumed that selecting “no action” is itself an action if such is feasible and potentially rational in the described decision situation. The nature of an action is open and it can refer to a certain quantity, but also to an arbitrarily complex procedure such as a computer programme in a certain programming language.

For each player $i \in N$ we denote her assigned **action set** by A_i . In principle, the action set A_i is arbitrary and can consist of any number of actions from which player i chooses. We emphasise that this assignment is a part of the modelling of the decision situation under consideration.

Players and their assigned actions are brought together in a structure that is referred to as a normal form pre-game. The pre-game omits a description of the consequences of the decisions made by the players. This is formalised in the next definition.

Definition 1.1 (Pre-games)

Let $N = \{1, \dots, n\}$ with $n \in \mathbb{N}$ be a finite set of players.

- (a) A **(normal form) pre-game** on N is a list $(A_i)_{i \in N} = (A_1, \dots, A_n)$ of action sets assigned to all the players in N , where A_i with $|A_i| \geq 2$ is the action set of player $i \in N$.
- (b) A pre-game $(A_i)_{i \in N}$ on N is **finite** if for every player $i \in N$ the associated action set A_i is a finite set.
- (c) In a pre-game $(A_i)_{i \in N}$ on N , an **profile** is an ordered list or action tuple $a = (a_1, \dots, a_n)$ with $a_i \in A_i$ an assigned action for player $i \in N$.
- (d) For a pre-game $(A_i)_{i \in N}$, we define the **profile space** as $\mathbf{A} = \prod_{i \in N} A_i = A_1 \times \dots \times A_n = \{(a_1, \dots, a_n) \mid a_i \in A_i \text{ for all } i \in N\}$ the Cartesian product of all players' action sets, being the set of all profiles in the pre-game.



The high level of abstraction of this pre-game model allows it to be applied to a wide variety of situations. A player may be an individual human being or any other decision-making entity like a government, a board of directors, the leadership of a revolutionary movement, or even a flower or an animal. The model places no restrictions on the set of actions available to a player, which may, for example, contain just a few elements or be a huge set containing complicated plans that cover a variety of contingencies.

The next example describes a number of situations that are represented as pre-games. We describe the decision makers as players and the actions that they can choose from.

Example 1.1 We consider a number of decision situations that can be represented as pre-games. The next examples indicate the wide variety of possibilities for representing decision situations as such.

The “rock-paper-scissors” game: Consider two persons with opposing views on a certain collective decision.

They could settle their dispute by playing the simple “rock-paper-scissors” game. This can be represented as a two-player game, with $N = \{1, 2\}$, each given the same decision set: $A_1 = A_2 = \{R, P, S\}$. Here, “R” stands for the action of selecting *Rock*; “P” stands for the action of selecting *Paper*, and; “S” stands for the action of selecting *Scissors*. Consequences of selection of these actions will be discussed later in Example 1.3.

A simple voting game: Consider a bowling club that holds a general meeting to determine whether their weekly bowling evening should move from a Wednesday to a Thursday. If there are $n = 34$ members, we can represent the vote on this proposed measure as a pre-game with 34 players, each having three actions at their disposal, namely to vote “YES”, to vote “NO”, or to “ABSTAIN”. Hence, the player set is $N = \{1, \dots, 34\}$ and for each player $i \in N$: $A_i = \{\text{YES}, \text{NO}, \text{ABSTAIN}\}$. Alternatively, we could use a numeric representation of these choices by letting $A_i = \{-1, 0, 1\}$, where $a_i = -1$ interprets as voting

“NO”, $a_i = 0$ interprets as voting “ABSTAIN”, and $a_i = 1$ interprets as voting “YES”.³

We do not discuss the results of these elections, since these depend on the exact rule used to determine a profile. Some of such rules and procedures are discussed later in Example 1.3.

Provision of a collective good: Consider a home owners association that wants to provide a playground to its members on the privately owned grounds in the association’s neighbourhood. Funds for this project will be raised through voluntary contributions from the households that are member of the association.

If there are $n \geq 2$ members of this association, we can model this situation as a pre-game with player set $N = \{1, \dots, n\}$ and action sets $A_i = [0, M_i] \subset \mathbb{R}_+$, for every player $i \in N$, where $M_i > 0$ denotes the available income is household $i \in N$. How the raised funds are exactly spent is omitted from the description of the decision situation as a pre-game. We refer to Example 1.3 for further discussion.

A Cournot duopoly: Consider a duopolistic market with two competing firms that produce an identical product. Both firms compete through determining their respective output quantities. Assuming a linear inverse demand function given by $P = \alpha - \beta Q$, where $\alpha, \beta > 0$ are demand parameters, $P \geq 0$ is the market price, and $Q \geq 0$ is the total quantity supplied in the market, we can represent this competitive market interaction as a pre-game. Indeed, the two firms are the sole decision makers, represented as player set $N = \{1, 2\}$. Each firm $i = 1, 2$ selects an output level $Q_i \in A_i = [0, \alpha]$. The total quantity provided in the market is now determined as $Q = Q_1 + Q_2$.

The examples above represent just a few cases of the variety of applications of this simple interactive decision model. ◆

This abstraction is at once a strength and a weakness: it is the abstraction that permits such wide applicability, but at this level of generality very few conclusions can be drawn about the outcome of a game. To derive interesting results one must be more specific.

The examples of the provision of a collective good as well as the Cournot duopoly as discussed above in Example 1.1 give rise to the definition of a certain class of pre-games, namely the ones with specific mathematical structures. In particular, it is widely accepted to consider games in which action sets have the standard Euclidean topological and ordinal structures. Hence, actions are themselves “vectors” in a Euclidean space and can be compared with each other through tools such as the Euclidean distance and the standard ordering of vectors in Euclidean spaces.

Euclidean pre-games The next definitions introduce the necessary mathematical tools to introduce this important class of pre-games.⁴



Mathematical notes Let $k \in \mathbb{N}$ be some natural number. The k -dimensional Euclidean vector space \mathbb{R}^k is the set of all k -dimensional vectors that is endowed with the topology generated by the standard Euclidean norm:

- $x \in \mathbb{R}^k$ is a vector $x = (x_1, \dots, x_k)$ where $x_m \in \mathbb{R}$ are real valued entries for all $m = 1, \dots, k$.
- The Euclidean norm of a vector $x \in \mathbb{R}^k$ is defined as $\|x\| = \sqrt{\sum_{m=1}^k x_m^2}$.
- The Euclidean distance between two vectors $x, y \in \mathbb{R}^k$ is defined by $d(x, y) = \|x - y\|$. We remark that $d: \mathbb{R}^k \times \mathbb{R}^k \rightarrow \mathbb{R}_+$ is a metric on \mathbb{R}^k that generates the standard Euclidean topology on this vector space.

³Other numerical representations of these actions are also possible. Actually, every triple of numbers $a, b, c \in \mathbb{R}$ could be used to describe these three actions as long as these three numbers are different.

⁴I refer to some standard textbooks on linear algebra, in particular linear vector spaces, by Axler (2020, 2024) and Strang (2021) for further details of the mathematics of topological vector spaces.

With regard to the ordering of the k -dimensional Euclidean space, we use standard notation and define for any pair of vectors $x, y \in \mathbb{R}^k$ the following relationships:

- $x \geq y$ if and only if $x_m \geq y_m$ for all $m = 1, \dots, k$;
- $x > y$ if and only if $x \geq y$ and $x_m > y_m$ for at least one $m \in \{1, \dots, k\}$, and;
- $x \gg y$ if and only if $x_m > y_m$ for all $m = 1, \dots, k$.

These relationships define the standard Euclidean ordering on the vector space \mathbb{R}^k . □

Using these notions, we now introduce two important topological notions that help us properly define a relevant class of pre-games on a player set. These notions are that of convergence and closedness of subsets in Euclidean spaces.



Mathematical notes Let \mathbb{R}^k be some k -dimensional Euclidean space.

- A **sequence** in \mathbb{R}^k is an ordered subset $\{x^m \in \mathbb{R}^k \mid m \in \mathbb{N}\}$ —usually denoted as $(x^m)_{m \in \mathbb{N}} \subset \mathbb{R}^k$.⁵
- The sequence $(x^m)_{m \in \mathbb{N}} \subset \mathbb{R}^k$ **converges** to a vector $x \in \mathbb{R}^k$ if it holds that for every $\varepsilon > 0$ there exists some $N \in \mathbb{N}$ such that $\|x - x^m\| < \varepsilon$ for every $m > N$. It is also said that x is the **limit** of the sequence $(x^m)_{m \in \mathbb{N}}$. If the sequence $(x^m)_{m \in \mathbb{N}}$ converges to $x \in \mathbb{R}^k$ we denote this as $x^m \rightarrow x$.
- A subset $S \subset \mathbb{R}^k$ is called **closed** in the Euclidean topology if for every convergent sequence $(x^m)_{m \in \mathbb{N}} \subset S$ with $x^m \rightarrow x$, it holds that $x \in S$. Hence, every limit of a convergent sequence in S is an element of S .
- A subset $S \subset \mathbb{R}^k$ is called **compact** in the Euclidean topology if S is closed and bounded in the sense that there exists some $B > 0$ such that $\|x\| < B$ for all $x \in S$.

These standard Euclidean topological concepts are also set out in Apostol (1974). □

These additional mathematical tools help us define the notion of a specific class of pre-games, which are founded on the topological structure of Euclidean spaces. These pre-games can be referred to as *Euclidean* to distinguish them from the general class of pre-games.

Definition 1.2 (Euclidean pre-games)

A pre-game $(A_i)_{i \in N}$ on the finite player set N is a **Euclidean pre-game** if for every player $i \in N$, the action set A_i is a closed subset of some finite dimensional Euclidean vector space, i.e., $A_i \subseteq \mathbb{R}^{k_i}$ for some $k_i \in \mathbb{N}$.

In a Euclidean pre-game the profile set \mathbf{A} is a closed subset of a finite dimensional Euclidean vector space as well. Indeed, the profile set $\mathbf{A} = \prod_{i \in N} A_i \subseteq \mathbb{R}^k$ is closed in the Euclidean topology on \mathbb{R}^k , where $k = \sum_{i \in N} k_i \in \mathbb{N}$. ♣

It is clear that the provision of a collective good pre-game and the Cournot pre-game as discussed in Example 1.1 are indeed Euclidean pre-games. Both represent decision situations in which players have a continuum of actions, clearly defined as real numbers. Therefore, the basic premise of a Euclidean pre-game is satisfied.

On the other hand, the class of Euclidean pre-games is much wider. Our discussion of pre-games with *finite* action sets for all players can also be understood as a Euclidean pre-game. In particular, this has been shown directly for the simple voting game in which the three actions “YES”, “NO”, and “ABSTAIN” are replaced by the numbers 1, -1 and 0, respectively.

This conversion can be applied to any pre-game with finite action sets for all players. This requires a straightforward transformation of the action sets into the real number line.

⁵In particular, we remark that the real number space \mathbb{R} itself is a Euclidean vector space endowed with the standard (absolute) norm $|\cdot|$ and the standard linear order on the real line.

Proposition 1.1

All finite pre-games are Euclidean.



Proof Consider a pre-game $(A_i)_{i \in N}$ on the player set N such that for every player $i \in N$ the assigned action set A_i is finite. In particular, let $A_i = \{a_i^1, \dots, a_i^{m_i}\}$, where $m_i \in \mathbb{N}$ is the number of actions available to player i . Now we convert A_i into the corresponding Euclidean action set $B_i = \{1, \dots, m_i\} \subset \mathbb{N}$. Hence, this shows that the finite action set is equivalent to some (finite) Euclidean set.

Let $\mathbf{A} = \prod_{i \in N} A_i$ be the finite profile set in the original pre-game and let $\mathbf{B} = \prod_{i \in N} B_i \subset \mathbb{R}^N$. Define the mapping $\chi: \mathbf{B} \rightarrow \mathbf{A}$ by $\chi(b) = (a_1^{b_1}, a_2^{b_2}, \dots, a_n^{b_n})$, which assigns to a numerical representation b the (unique) corresponding profile $\chi(b) \in \mathbf{A}$ in the original game. It is clear that $\chi(\mathbf{B}) = \mathbf{A}$ and $\chi^{-1}(\mathbf{A}) = \mathbf{B}$, showing that χ is indeed a one-to-one bijection between \mathbf{A} and \mathbf{B} .

This shows that the pre-game $(B_i)_{i \in N}$ is a transformation of the original pre-game $(A_i)_{i \in N}$. Since $(B_i)_{i \in N}$ is a Euclidean pre-game, it follows that the original (finite) pre-game $(A_i)_{i \in N}$ is mathematically equivalent to a Euclidean game, showing the asserted property. \square

Proposition 1.1 implies that every pre-game discussed in Example 1.1 is either Euclidean or may be represented as such. A non-Euclidean pre-game must therefore have infinite action sets. The next example illustrates two such pre-games.

Example 1.2 Let $N = \{1, 2\}$ be a player set with two players. We consider the following pre-games for two players.

Simplified bargaining Consider a multi-round bargaining situation where two players negotiate about splitting a cake, using “demands” as the main communicating device. Assuming the size of the cake is one, the bargaining process consists of an indeterminate number of rounds. In each round $t \in \mathbb{N}$, player 1 puts forward a demand $\alpha_t \in [0, 1]$, while player 2 puts forward $\beta_t \in [0, 1]$. The bargaining terminates in period t if and only if $\alpha_t + \beta_t \leq 1$, i.e., the bargaining concludes when it is feasible to pay out the demands of both players.

To structure this as a normal form pre-game, we are required to introduce actions that can be executed throughout the bargaining process. For that purpose we introduce actions that are infinite tuples of demands. Hence, for player 1 an action would be given by $a_1 = (\alpha_t)_{t \in \mathbb{N}} \in A_1 = [0, 1]^{\mathbb{N}}$ and for player 2 an action can be modelled as $a_2 = (\beta_t)_{t \in \mathbb{N}} \in A_2 = [0, 1]^{\mathbb{N}}$.

Clearly, this is not a Euclidean pre-game, since the action sets $A_1 = A_2 = [0, 1]^{\mathbb{N}}$ are infinitely dimensional vector spaces.

Advanced financial duopoly During the past decade, financial trade technology has made significant advances with the introduction of trade algorithms as well as “artificial intelligence” (AI). These technologies help financial brokers to trade in financial capital markets at the speed of light with nano-second trade periods. This has resulted in anomalous trade patterns in these markets. We can give the following description of a pre-game model of such a market with two traders.

Consider the players to be two brokers in a system of financial markets with multiple assets $\ell = 1, \dots, m$, where m is sufficiently large. A trade is now a list of net trades $\tau = (\tau_1, \dots, \tau_m) \in \mathbb{R}^m$. Trades are put to the market system by both traders and are reconciled through a well-defined and well-described demand and supply reconciliation system.⁶ An action is now a computer programme or Turing machine

⁶I omit here a full description of such a market mechanism, but it can be formulated using standard expressions from market economics.

in computer language “L” that uses past trades as inputs and determines a trade. An action set is now the collection of all theoretically possible actions—or computer programmes in “L”—as described.

Clearly, the action set consists of arbitrarily complex computer programmes that cannot properly be expressed in numerical terms, let alone as finite Euclidean vectors. Therefore, this pre-game is also non-Euclidean. ◆

1.2.2 Normal form games: Payoff structures

In standard game theory, the specification is completed by adding a payoff structure to a pre-game to create a normal form or strategic game. A payoff structure assigns to every profile in the pre-game a certain numerical value for every individual player in the pre-game. Hence, a payoff structure evaluates every profile of the pre-game through the assignment of a (real) numerical value. The interpretation of this evaluation is left open in the definition of a normal form game. Indeed, the assigned values or “payoffs” can be interpreted cardinally as well as ordinally.

Payoffs as cardinals means that they describe real values such as monetary amounts or profiles on a numerical scale of measurement such as temperatures, weights, or lengths. This implies that computations and manipulations can be performed on these cardinal payoffs, appropriate according to the interpretation given to these cardinal values.

On the other hand, ordinal payoffs are interpreted as tools in ranking profiles: a profile with a higher payoff is preferred over a profile with a lower payoff. This implies that no computations and manipulations can be performed with these payoffs; they are just numerical values that represent ordinal rankings.

The next definition introduces a formalisation of a payoff structure on a pre-game and brings the pre-game together with a payoff structure into a “game”.

Definition 1.3 (Normal form games)

Let $N = \{1, \dots, n\}$ be a finite set of players.

- (a) A **payoff structure** on a pre-game $(A_i)_{i \in N}$ on player set N as defined in Definition 1.1 is a list $(\pi_i)_{i \in N}$ with for every player $i \in N$, an assigned payoff function $\pi_i: \mathbf{A} \rightarrow \mathbb{R}$ that links to every profile $a = (a_1, \dots, a_n) \in \mathbf{A} = \prod_{i \in N} A_i$ a payoff $\pi_i(a) \in \mathbb{R}$ to player $i \in N$.

We can also denote a payoff structure on a pre-game $(A_i)_{i \in N}$ as an n -dimensional mapping $\pi = (\pi_1, \dots, \pi_n): \mathbf{A} \rightarrow \mathbb{R}^N$.

- (b) A **(normal form) game** on player set N is a combination of a pre-game $(A_i)_{i \in N}$ and a payoff structure $(\pi_i)_{i \in N}$ denoted as $(A_i, \pi_i)_{i \in N}$.

Throughout we use the symbolic notation $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ to denote a normal form game on player set N . A game Γ on player set N is also known as an “ n -player game”.

- (c) The **space of n -player games** is defined as

$$\mathbb{G}^N = \{ \Gamma \mid \Gamma = \langle N, \mathbf{A}, \pi \rangle \text{ is a normal form game on } N \} \quad (1.2)$$

The **universal game space** is the set of all finite player games defined by

$$\mathbb{G} = \bigcup_{N \in \mathcal{N}} \mathbb{G}^N \quad (1.3)$$

with $\mathcal{N} = \{N \mid N = \{1, \dots, n\} \text{ for some } n \in \mathbb{N}\}$.



The pre-game captures only the abstract structure of social interaction between the players; the payoff structure completes the description by introducing a subjective, individualistic assessment of its profiles. Payoffs evaluate the *consequences* of the decision process and impose on the players a ranking of profiles that gives direction to purposeful or “rational” behaviour in the game.

Normal form games can be enhanced with elements that describe outside events, such as changes in the “state of the world” through random events. Decisions are then taken before the state of the world is realised, and payoffs are replaced by expected payoffs.

Example 1.3 We consider the pre-games introduced in Example 1.1.

The “rock-paper-scissors” game: Recalling the RPS pre-game described in Example 1.1, we have to introduce a payoff structure that reflects the standard ordinal ranking of actions: R beats S; S beats P, and; P beats R. All equal choices are a draw.

The payoff structure $\pi : \{R, P, S\} \times \{R, P, S\} \rightarrow \mathbb{R}^2$ should reflect this ranking. Hence, π represents a purely ordinal ranking and should not be construed as being cardinal.

We can describe π as follows:

$$\begin{aligned} \pi(R, R) &= \pi(S, S) = \pi(P, P) = (0, 0) \\ \pi(R, S) &= \pi(S, P) = \pi(P, R) = (1, -1) \\ \pi(S, R) &= \pi(P, S) = \pi(R, P) = (-1, 1) \end{aligned}$$

It is clear that the payoff structure of this game is based on the assignment of only three numerical values, 0, 1 and -1 , which are used to rank profiles.

The payoff structure can also be represented in the form of a table. The vertical dimension is the choice of the first player, while the horizontal dimension represents the choice of the second player. This results in a two-dimensional payoff grid:⁷

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

A game with two players and a finite number of actions for each of these players is also known as a *matrix game*, since any payoff structure can be represented as a table or matrix as depicted for the Rock-Paper-Scissors game.

⁷In this matrix, we use the convention that player 1 is the row player, selecting the *R*-, *P*- or *S*-row in the matrix, while player 2 is the column player, selecting the *R*-, *P*- or *S*-column in the matrix. Each entry in this matrix represents a unique profile in this game from (*R*, *R*) through (*S*, *S*). In the corresponding field, we put the payoffs of both players, using the rule that the first payoff is that of the first player (the “row player”) and the second payoff is that of the second player (the “column player”). Hence, $\pi(R, R) = (0, 0)$ is entered in the upper left corner field corresponding to the profile (*R*, *R*). All other payoffs are linked to profiles in a similar fashion.

A simple voting game: In the voting game described in Example 1.1, decisions are made based on majority for a certain decision. We consider here two different methods to put this in practice, a simple and an absolute majority rule.⁸

Simple majority rule — Under a simple majority rule, we can use numerical expressions of votes ($a_i \in \{-1, 0, 1\}$ for all $i \in \{1, \dots, 34\}$, as used in Example 1.1) to introduce a mathematical expression. Indeed, the group goes ahead with the move of the bowling evening to Thursday if and only if

$$\sum_{i=1}^{34} a_i \geq 1.$$

This rule expresses that there is majority of YES-votes in comparison with NO-votes. Note that ABSTAIN-votes are not really counted. This translates to a payoff structure that expresses this decision in a payoff, for example, as follows:

$$\pi_i(a) = \begin{cases} \alpha_i & \text{if } \sum_{i=1}^{34} a_i \geq 1 \\ \beta_i & \text{if } \sum_{i=1}^{34} a_i \leq 0 \end{cases}$$

where $(\alpha_i, \beta_i) \in \mathbb{R}^2$ denote the individual payoffs assigned by player $i \in N$ to the two possible collective decisions.

Absolute majority rule — Under the absolute majority rule, the bowling club goes ahead with moving the bowling evening to Thursday if and only if an absolute majority of at least 18 YES-votes is in favour of the building of the playground. Hence, this leads in turn to a payoff structure given by

$$\pi_i(a) = \begin{cases} \alpha_i & \text{if } \#\{i \in N \mid a_i = 1\} \geq 18 \\ \beta_i & \text{if } \#\{i \in N \mid a_i = 1\} \leq 17 \end{cases}$$

where $(\alpha_i, \beta_i) \in \mathbb{R}^2$ again denote the individual payoffs assigned by player $i \in N$ to the two possible collective decisions.

Provision of a collective good: The most plausible decision rule for the home association is to build the playground if and only if the collectively raised funds cover its construction cost. Hence, assume the cost of building the playground is given by $0 < C \leq \sum_{i \in N} M_i$. Then, the association constructs a playground if and only if

$$\sum_{i \in N} a_i \geq C.$$

Assuming that players do not have any opportunity costs from using their budget M_i , this can be translated into a payoff structure given by

$$\pi_i(a) = \begin{cases} \alpha_i - a_i & \text{if } \sum_{i \in N} a_i \geq C \\ \beta_i & \text{if } \sum_{i \in N} a_i < C \end{cases}$$

where $(\alpha_i, \beta_i) \in \mathbb{R}^2$ again denote the individual gross benefits assigned by player $i \in N$ to the two

⁸It is clear that there are numerous alternative ways to make a decision about the collective project, each can be modelled through systems of inequalities as described for these two decision rules.

possible collective decisions of the playground being built or not.

A Cournot duopoly: A standard hypothesis in the Cournot duopoly model is that both firms are profit maximisers. If the two firms produce Q_1 and Q_2 , respectively, then the resulting market price can be computed as

$$P(Q_1, Q_2) = \alpha - \beta(Q_1 + Q_2)$$

and the resulting profits are determined as

$$\pi_1(Q_1, Q_2) = P(Q_1, Q_2) \cdot Q_1 - c_1(Q_1)$$

$$\pi_2(Q_1, Q_2) = P(Q_1, Q_2) \cdot Q_2 - c_2(Q_2)$$

where $c_1, c_2: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ are the cost functions of the two respective firms in this duopoly.

These examples of normal form games show how payoff structures give direction to and accommodate purposeful behaviour of the players in their decision making processes. ◆

Normal form games form a powerful class of descriptive tools for social interactive decision making, accommodating the introduction of purposeful or “rational” behaviour. I take this up at length in the next section.

I close this discussion with an extensive quotation from Osborne and Rubinstein (1994) that provides a concise assessment of the purpose of the normal form.⁹

Comments on interpretation

A common interpretation of a strategic game is that it is a model of an event that occurs only once; each player knows the details of the game and the fact that all the players are “rational” [...], and the players choose their actions simultaneously and independently. Under this interpretation each player is unaware, when choosing his action, of the choices being made by the other players; there is no information (except the primitives of the model) on which a player can base his expectation of the other players’ behavior.

Another interpretation, which we adopt through most of this book, is that a player can form his expectation of the other players’ behavior on the basis of information about the way that the game or a similar game was played in the past [...]. A sequence of plays of the game can be modeled by a strategic game only if there are no strategic links between the plays. That is, an individual who plays the game many times must be concerned only with his instantaneous payoff and ignore the effects of his current action on the other players’ future behavior. In this interpretation it is thus appropriate to model a situation as a strategic game only in the absence of an intertemporal strategic link between occurrences of the interaction. [...]

When referring to the actions of the players in a strategic game as “simultaneous” we do not necessarily mean that these actions are taken at the same point in time. One situation that can be modeled as a strategic game is the following. The players are at different locations, in front of terminals. First the players’ possible actions and payoffs are described publicly (so that they are common knowledge among the players). Then each player chooses an action by sending a message to a central computer; the players are informed of their payoffs when all the messages have been received. However,

⁹I also refer to Hargreaves-Heap and Varoufakis (2004) for elaborate, critical assessments of game theory and the various mathematical game forms used to describe social interactive decision making.

the model of a strategic game is much more widely applicable than this example suggests. For a situation to be modeled as a strategic game it is important only that the players make decisions independently, no player being informed of the choice of any other player prior to making his own decision.

Osborne and Rubinstein (1994), Section 2.1.2, pages 13–14.

Euclidean games Having introduced the notion of a Euclidean pre-game in the previous section, we may extend this particular definition to the realm of games as well. For this we need to introduce a few more auxiliary mathematical concepts.



Mathematical notes Let $(A_i)_{i \in N}$ be some Euclidean pre-game on the player set N , where all action sets A_i are closed subsets of some finite dimensional Euclidean vector space. The generated profile set $\mathbf{A} = \prod_{i \in N} A_i$ is, therefore, also a closed subset of a finite dimensional Euclidean vector space.

- For a player $i \in N$ the payoff function $\pi_i: \mathbf{A} \rightarrow \mathbb{R}$ is **continuous at profile** $a \in \mathbf{A}$ if for every convergent sequence $a^n \rightarrow a$ in \mathbf{A} we have that $\pi_i(a^n) \rightarrow \pi_i(a)$.
- For a player $i \in N$ the payoff function $\pi_i: \mathbf{A} \rightarrow \mathbb{R}$ is **continuous** if π_i is continuous at every profile $a \in \mathbf{A}$.

If the payoff functions π_i are continuous on \mathbf{A} for all players $i \in N$, then the (n -dimensional) payoff function $\pi: \mathbf{A} \rightarrow \mathbb{R}^N$ is continuous as well. \square

We use these mathematical notions to introduce the notion of a Euclidean normal form game, usually denoted simply as a Euclidean game. This is complemented with the notion of a compact game, in which all action sets and the resulting profile set are compact in the Euclidean topology.

Definition 1.4

A normal form game $\Gamma = \langle N, \mathbf{A}, \pi \rangle \in \mathbb{G}^N$ on the player set N is **Euclidean** if the corresponding pre-game $(A_i)_{i \in N}$ is Euclidean and for every player $i \in N$ the payoff function $\pi_i: \mathbf{A} \rightarrow \mathbb{R}$ is continuous on \mathbf{A} . The set of Euclidean games on N is denoted by

$$\mathbb{G}_E^N = \{ \Gamma \in \mathbb{G}^N \mid \Gamma = \langle N, \mathbf{A}, \pi \rangle \text{ is Euclidean} \}. \quad (1.4)$$

A Euclidean game $\Gamma \in \mathbb{G}_E^N$ is **compact** if for every player $i \in N$ the action set $A_i \subset \mathbb{R}^{k_i}$ is compact in the Euclidean topology on \mathbb{R}^{k_i} . The set of compact (Euclidean) games on N is denoted by

$$\mathbb{G}_C^N = \{ \Gamma \in \mathbb{G}_E^N \mid \Gamma = \langle N, \mathbf{A}, \pi \rangle \text{ is compact} \}. \quad (1.5)$$

Clearly, $\mathbb{G}_C^N \subset \mathbb{G}_E^N \subset \mathbb{G}^N$.



1.2.3 Nash equilibrium

The most common approach to defining “rational behaviour” in a normal form game is the notion of Nash equilibrium, seminaly introduced by Nash (1950). A Nash equilibrium represents a *steady state* of game play at which no individual player can improve her payoff through a unilateral deviation. It is the baseline profile that emerges from rational action selection: it assumes intelligent decision making on the part of every player, aimed at optimising her own payoff given the payoff structure of the game.

To formalise this behavioural norm, we introduce some additional notation. Let (N, \mathbf{A}) be a pre-game and let $a = (a_1, \dots, a_n) \in \mathbf{A} = \prod_{i \in N} A_i$ be some profile in that pre-game. Then for player $i \in N$ we denote by a_{-i} the list of all players' actions except the action of player $i \in N$. Hence,

$$a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n) \in \mathbf{A}_{-i} = \prod_{j \neq i} A_j.$$

This implies that for any $i \in N$: $a = (a_i, a_{-i}) \in \mathbf{A}$. this means that player $i \in N$ deviating from profile $a \in \mathbf{A}$ to move to a modified profile in the pre-game can be represented by $(b_i, a_{-i}) \in \mathbf{A}$, where $b_i \in A_i$ is the selected alternative action by player i .

We can now formalise the notion of Nash equilibrium.

Definition 1.5 (Nash equilibrium)

Let $\Gamma = \langle N, \mathbf{A}, \pi \rangle \in \mathbb{G}^N$ be a normal form game. Then the profile $a^* \in \mathbf{A}$ is a **Nash equilibrium** in Γ if for every player $i \in N$ it holds that

$$\pi_i(a^*) \geq \pi_i(b_i, a_{-i}^*) \tag{1.6}$$

for every action $b_i \in A_i$.

The set of Nash equilibria in the game $\Gamma \in \mathbb{G}^N$ is denoted by $\text{NE}(\Gamma) \subset \mathbf{A}$.

The profile $a^* \in \mathbf{A}$ is a **strict Nash equilibrium** in Γ if for every player $i \in N$ it holds that

$$\pi_i(a^*) > \pi_i(b_i, a_{-i}^*) \tag{1.7}$$

for every alternative action $b_i \in A_i$ with $b_i \neq a_i^*$.



Nash equilibrium thus describes a stable profile of a game in the sense that no individual player has any incentive to deviate from her Nash equilibrium action. This can be illustrated by considering a few simple games and determining their Nash equilibria.

Example 1.4 Let $N = \{1, 2\}$ be a two-player set. We consider the following two-player games to illustrate the definition of the Nash equilibrium concept.

- (i) One of the most well-known games is that of the ‘‘Prisoners’ Dilemma’’. It refers to a decision situation with two prisoners accused of a serious crime dealing with a prosecutor. If both of them do not talk and are uncooperative, the prosecutor can only show a minor offence, resulting in a lenient sentence for both of them. If one of the two prisoners confesses and the other prisoner remains silent, the prosecutor will set the confessing prisoner free, while the other, silent prisoner is given a lengthy sentence. Finally, if both players talk and confess, both will receive substantial sentences, though less than the lengthy sentence.

This can be described by the same action set for both players, $A_1 = A_2 = \{C, D\}$ —referred to as a cooperative action ‘‘C’’ and a defecting action ‘‘D’’.¹⁰ The assigned prison sentences can be described in the following matrix:

¹⁰Here the cooperation considered is that with the other prisoner, not the prosecutor. Hence, ‘‘C’’ refers to solidarity with the fellow prisoner and ‘‘D’’ refers to the abandonment of the other prisoner.

	C	D
C	-1, -1	-10, 0
D	0, -10	-5, -5

In this simple game (D, D) is the unique Nash equilibrium and that this Nash equilibrium is strict. Indeed, $\pi_1(D, D) = -5 > \pi_1(C, D) = -10$ as well as $\pi_2(D, D) = -5 > \pi_2(D, C) = -10$. Note that (C, C) is no Nash equilibrium since $\pi_1(C, C) = -1 < \pi_1(D, C) = 0$. Similarly, (D, C) and (C, D) can be ruled out as Nash equilibria.

- (ii) Consider the “Rock-Paper-Scissors” game as discussed in Example 1.3. In this matrix game, there does not exist any Nash equilibrium. This can be checked for every potential profile $(a_1, a_2) \in \{R, P, S\}^2$. For example, (R, R) is not a Nash equilibrium since $\pi_1(P, R) = \pi_2(R, P) = 1 > \pi_1(R, R) = \pi_2(R, R) = 0$, hence both players prefer to deviate from selecting R by selecting P . Furthermore, (P, R) is not a Nash equilibrium either, since $\pi_2(P, S) = 1 > \pi_2(P, R) = -1$, indicating that player 2 would deviate from R by selecting S instead.
- (iii) Next, consider the simple voting game under the simple majority rule as discussed in Examples 1.1 and 1.3. In principle there exist many Nash equilibria in this game, depending on the exact utilities (α_i, β_i) for all $i \in N$. For an arbitrary value distribution over the 34 players, we determine a straightforward voting strategy given by

$$\hat{a}_i = \begin{cases} 1 & \text{if } \alpha_i \geq \beta_i \\ -1 & \text{if } \alpha_i < \beta_i \end{cases}.$$

This straightforward voting strategy $\hat{a} \in \mathbf{A}$ results in the move of the bowling evening to Thursday if and only if $\#\{i \in N \mid \alpha_i \geq \beta_i\} > \#\{j \in N \mid \alpha_j < \beta_j\}$, implying that \hat{a} is indeed a Nash equilibrium of this game. Either, deviating from this voting strategy has no effect, or deviating would not improve the profile of the vote for the individual player.

Assuming $\#\{i \in N \mid \alpha_i \geq \beta_i\} > \#\{j \in N \mid \alpha_j < \beta_j\}$, any profile $\tilde{a} \in \mathbf{A}$ in which any player $j \in N$ with $\alpha_j < \beta_j$ switches to $\tilde{a}_j = 1$ instead of $\hat{a}_j = -1$ will change the profile of the vote. Hence, \tilde{a} is also a Nash equilibrium. This shows that this simple voting game admits in principle many Nash equilibria.¹¹

- (iv) Finally, consider the provision of a public good as discussed in Examples 1.1 and 1.3. Let $N_1 = \{i \in N \mid \alpha_i > \beta_i\}$ be the set of players that recognise the benefits of the playground. The complement $N_2 = \{j \in N \mid \alpha_j \leq \beta_j\}$ are the players that do not think building a playground is strictly beneficial or desired. Note that N_1 and N_2 form a partition of N .

Note that, given the stated payoff function π_i in Example 1.3, each player $i \in N_1$ will contribute at most $\bar{a}_i = \min\{M_i, \alpha_i - \beta_i\} > 0$ to the building project. Hence, the building project is actually executed if and only if $\sum_{i \in N_1} \bar{a}_i \geq C$. In particular note that any financing strategy $0 \leq \hat{a}_i \leq \bar{a}_i$ with $\sum_{i \in N_1} \hat{a}_i = C$ is a Nash equilibrium of this public good provision game.

This shows that this game has infinitely many Nash equilibria if and only if $\sum_{i \in N_1} \bar{a}_i > C$.

¹¹The reader is invited to carefully check these claims and the corresponding Nash equilibrium profiles.

These examples show that there are games with no Nash equilibria, one equilibrium, a finite number of Nash equilibria, or infinitely many Nash equilibria. \blacklozenge

The Nash equilibrium conception refers to that all players optimise their chosen action to maximise their payoffs *given the actions chosen by the other players*. In this regard, these players “respond” optimally to what other players have chosen. It reflects, therefore, *best response rationality*. Best response rationality is a particular form of “intelligent” behaviour.

The next definition explores best response rationality. It introduces the mathematical tools to describe the selection of optimal responses to other players’ actions.

Definition 1.6 (Best responses)

Let $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ be a normal form game and let $i \in N$ be some player in Γ .

- (a) Let $a_{-i} \in \mathbf{A}_{-i}$. Then $\hat{a}_i \in A_i$ is a **best response** to a_{-i} if $\pi_i(\hat{a}_i, a_{-i}) \geq \pi_i(b_i, a_{-i})$ for every $b_i \in A_i$, or

$$\pi_i(\hat{a}_i, a_{-i}) = \max_{b_i \in A_i} \pi_i(b_i, a_{-i}) \quad (1.8)$$

- (b) For $a_{-i} \in \mathbf{A}_{-i}$, the **best response set** is defined by

$$B_i(a_{-i}) = \arg \max_{b_i \in A_i} \pi_i(b_i, a_{-i}) \quad (1.9)$$

This introduces the **best response correspondence** for player $i \in N$ by $B_i: \mathbf{A}_{-i} \rightarrow 2^{A_i}$. This summarises as the best response correspondence $B = (B_1, \dots, B_n): \mathbf{A} \rightarrow 2^{\mathbf{A}}$ for the game Γ . \clubsuit

These definitions articulate a particular form of intelligent decision making, namely *best response rationality*. A player who selects a best response to the actions of the others chooses an action that optimises her payoff function, given what the others have selected.

If all players exhibit best response rationality, the game reaches a stable state at which every player’s action is a best response to the remaining profile. Formally, $a^* \in \mathbf{A}$ is such a stable state if for every $i \in N: a_i^* \in B_i(a_{-i}^*)$, or equivalently $a^* \in B(a^*)$. The stable state is therefore a *fixed point* of the best response correspondence.

In a strict Nash equilibrium this relationship is sharper: each player selects the unique best action in response to the others, and $a^* \in \mathbf{A}$ is a strict Nash equilibrium if and only if $B(a^*) = \{a^*\}$. Any unilateral deviation is therefore strictly costly, making a strict Nash equilibrium a stronger form of stable state than a regular Nash equilibrium.

The next theorem states that the set of stable states under best response rationality is exactly the set of Nash equilibrium of the game. We state this result without a proof, which is rather straightforward.

Theorem 1.1 (Equivalent definition of Nash equilibrium)

Let $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ be a normal form game. Then:

- (a) $a^* \in \mathbf{A}$ is a Nash equilibrium in Γ if and only if a^* is a fixed point of the best response correspondence B , i.e., $a^* \in B(a^*)$.

(b) $a^* \in \mathbf{A}$ is a strict Nash equilibrium in Γ if and only if $B(a^*) = \{a^*\}$. ♥

Best response rationality can be used quite easily to compute Nash equilibria of quite complex games. In fact, it requires us to search for a fixed point of the best response correspondence of that particular game.

Example 1.5 Again let $N = \{1, 2\}$ be a two-player set. We consider some games that illustrate how Nash equilibrium can be found by computing fixed points of the best response correspondence.

- (i) Consider the Cournot duopoly described in Example 1.3. We can write the resulting payoff functions as

$$\begin{aligned}\pi_1(Q_1, Q_2) &= \alpha Q_1 - \beta Q_1(Q_1 + Q_2) - c_1(Q_1) \\ \pi_2(Q_1, Q_2) &= \alpha Q_2 - \beta Q_2(Q_1 + Q_2) - c_2(Q_2)\end{aligned}$$

To determine the best responses of the two firms in this game, we note that these payoff functions are concave if the cost functions are weakly convex. Assuming this, a stationary points of these two payoff functions exactly correspond to the optimal actions. Hence, we compute all those profiles (Q_1, Q_2) such that $\pi'_1(Q_1, Q_2) = \pi'_2(Q_1, Q_2) = 0$.

As an example we let the cost functions for both firms be linear and given by $c_1(Q_1) = \gamma_1 Q_1$ and $c_2(Q_2) = \gamma_2 Q_2$ satisfying the feasibility conditions $\gamma_1, \gamma_2 \in [0, \alpha]$. We can determine the stationary points for this particular case as described. This results in

$$\begin{aligned}\pi'_1(Q_1, Q_2) &= \alpha - \gamma_1 - \beta(2Q_1 + Q_2) \equiv 0 \\ \pi'_2(Q_1, Q_2) &= \alpha - \gamma_2 - \beta(Q_1 + 2Q_2) \equiv 0\end{aligned}$$

or

$$Q_1 = B_1(Q_2) = \frac{\alpha - \gamma_1 - \beta Q_2}{2\beta} = \frac{\alpha - \gamma_1}{2\beta} - \frac{1}{2}Q_2 \quad (1.10)$$

$$Q_2 = B_2(Q_1) = \frac{\alpha - \gamma_2 - \beta Q_1}{2\beta} = \frac{\alpha - \gamma_2}{2\beta} - \frac{1}{2}Q_1 \quad (1.11)$$

First, we note that this system of two equations describes the complete best response correspondence, or in this case the *best response function* on \mathbb{R}_+^2 . We can now solve the system of these two equations (1.10) to find the fixed point of the best response function. This results in the Nash equilibrium¹² proper given by

$$Q_1^* = \frac{\alpha - 2\gamma_1 + \gamma_2}{3\beta} \quad \text{and} \quad Q_2^* = \frac{\alpha + \gamma_1 - 2\gamma_2}{3\beta}$$

resulting into a total produced output in the market of $Q^* = Q_1^* + Q_2^* = \frac{2\alpha - \gamma_1 - \gamma_2}{3\beta} \geq 0$ and an equilibrium market price $P^* = P(Q^*) = \frac{1}{3}(\alpha + \gamma_1 + \gamma_2) \leq \alpha$.

¹²Due to the historical origins of this market game and the fact that Cournot (1838) already determined and discussed this equilibrium, we may also refer to this as the Cournot-Nash equilibrium in a duopoly.

- (ii) Next, again consider the provision of a public good discussed in Examples 1.1, 1.3 and 1.4. As before, consider $N_1 = \{i \in N \mid \alpha_i > \beta_i\}$ as the set of players that recognise the benefits of the playground. The maximal contribution of player $i \in N_1$ was determined as $\min\{M_i, \alpha_i - \beta_i\}$. As a consequence, we already established that the building project is executed if and only if $\sum_{i \in N_1} \min\{M_i, \alpha_i - \beta_i\} \geq C$. We can express this also through determining the best responses of the players in N_1 . It can be established that the best response of player $i \in N_1$ with regard to $a_{-i} \in \mathbf{A}_{-i} = \prod_{j \neq i} A_j$ can be determined as

$$B_i(a_{-i}) = \begin{cases} \max\{0, C - \sum_{j \neq i} a_j\} & \text{if } \sum_{j \neq i} a_j \geq C - \min\{M_i, \alpha_i - \beta_i\} \\ 0 & \text{if } \sum_{j \neq i} a_j < C - \min\{M_i, \alpha_i - \beta_i\} \end{cases}$$

where we recall that $\min\{M_i, \alpha_i - \beta_i\}$ is the maximal amount that player $i \in N_1$ is willing to contribute to the collective good, in this case the construction of a playground. Note that indeed $B_i(a_{-i}) \leq \min\{M_i, \alpha_i - \beta_i\}$ for every $i \in N_1$ as described above.

- (iii) *Coordination games* form a specific class of normal form games in which players coordinate their actions to arrive at optimal profiles. The classical example is the coordination on a date. Two persons can go to a James Bond movie or otherwise go to a Tom Cruise movie, like some “Mission: Impossible” incarnation. One person prefers James Bond over a Tom Cruise movie, but prefers either over not meeting for a date. The other person has exactly opposite preferences, preferring Tom Cruise over James Bond, but agreeing that not having a date is worse than either. This can be described as a matrix game with both players have two strategies.

If one player is the row player, while the other player is the column player, this results in the following matrix representation:

	JB	TC
JB	2, 1	0, 0
TC	0, 0	1, 2

From a best response perspective, JB is a best response to JB for both players, making (JB, JB) a Nash equilibrium. Similarly, TC is a best response to TC for both players, implying that (TC, TC) another Nash equilibrium.

The obvious problem in this coordination game is actually on which Nash equilibrium to coordinate. Both of these Nash equilibria are actually socially optimal in the sense that there is no alternative profile that makes both of them better off. It is impossible to advice both players on which equilibrium to coordinate.

- (iv) Next, we look again at the coordination game in (ii), but for two different players. Now we assume that both players have exactly the same preferences over the two alternatives: Both prefer going to a James Bond movie over seeing a Tom Cruise feature. This can be represented by a matrix game described by

	JB	TC
JB	2, 2	0, 0
TC	0, 0	1, 1

Again (JB, JB) and (TC, TC) are the two Nash equilibria in this modified game. Now, however, the equilibrium (JB, JB) results in uniformly higher payoffs for both players in comparison with the alternative equilibrium (TC, TC). Hence, (JB, JB) *Pareto dominates* (TC, TC). This could function as a determinant in guiding both players in their decision making and they can more easily coordinate their actions on the Pareto optimal equilibrium (JB, JB).

These four examples illustrate the use of best response rationality, not only in computing the resulting equilibria, but also in the more detailed analysis of the behaviour of the players in the game. ♦

Interpretations of the Nash equilibrium concept The Nash equilibrium concept is *the quintessential equilibrium conception* in game theory and is used in many applications of game theory to related fields such as computer science, biology, and economics. The reasons that the Nash equilibrium conception plays such a crucial role is that it has some important properties and interpretations. This is explored in many publications and textbooks.¹³ I summarise some arguments about the usefulness of the Nash equilibrium concept here as well.

The foremost reason that Nash equilibrium is such a useful conception is that it represents a natural outcome of self-motivated behaviour under ignorance—what can be termed *behavioural stability*. Given that an individual player does not really know the reasoning process followed by other players in a game, she most naturally will take the observed behaviour of the other players as given. Therefore, it is most natural to select the action that optimises the outcome for every player *taking the actions of the other players as given and immutable*. Hence, the Nash equilibrium profile becomes a naturally “expected” outcome of the game and it is founded in all players’ ignorance about their respective behaviour. This interpretation allows the Nash equilibrium to be interpreted as a meta-solution: The stability property is one that we would like every natural and reasonable solution conception to exhibit. The Nash equilibrium concept can, therefore, be interpreted as a minimal behavioural benchmark. It is the most pure and natural form of “selfish” intelligent behaviour, namely founded in the logic of selecting the optimal action if all other players’ actions are taken as given.

A second interpretation of the Nash equilibrium concept is that of a *self-confirming agreement*, through the related notion of a self-confirming equilibrium as put forward by Fudenberg and Levine (1993). Here, an “agreement” is interpreted as a non-binding settlement between the players. No player will deviate from the proposed agreement because there is no possibility to benefit from such a deviation. This is particularly convincing in the context of the class of coordination games. In this regard, “self-confirming” can also be re-stated as “self-fulfilling”.

A third interpretation is that of a *normative recommendation*. One can consider a strategic form game also from a normative point of view. Suppose that there is an impartial arbitrator or regulator that oversees the interactions represented by the game. In this case, the arbitrator has to propose to all decision makers (the players) how they should behave. One would expect the arbitrator’s proposal to be a Nash equilibrium, i.e., an outcome that cannot be undermined by opposition from individual decision makers. In a Nash equilibrium this is indeed not the case, since individual players have no benefit from deviation from the suggested action. The problem of arbitration arises if there are multiple Nash equilibria in the game under consideration. In that case, secondary considerations enter the debate for selecting a recommended profile. The Nash equilibrium conditions become a minimal requirement, rather than the sole consideration for selecting a certain profile. From

¹³For example, I refer here to the discussions in Osborne and Rubinstein (1994, Chapters 2–5), Maschler, Solan, and Zamir (2013, Section 4.9) and Bonanno (2018, Section 2.6). For a philosophical discussion and treatment, I also refer to Weirich (1998).

the discussion above it should be clear that the Nash equilibrium concept is a natural outcome that functions as a *minimal* requirement for any consideration of play in a strategic form game. There has arisen a wide-ranging literature on equilibrium conceptions with discussions on what rationality or intelligent behaviour exactly entails in the context of strategic form games. These contributions make clear that the Nash equilibrium concept is a benchmark, but that other considerations can be important as well. Hence, even if a game has a unique Nash equilibrium, it might not be the “best” or “recommended” way to select actions in that game. (I refer here to the Prisoners’ Dilemma game discussed in Example 1.4.) This is besides the obvious issues that arise about equilibrium in games that have no Nash equilibrium or have multiple Nash equilibria. In those games, other criteria—besides Nash’s best response rationality—play a critical role in assessing how players ought to behave or actually will behave.

Properties of Nash equilibria If a normal form game is Euclidean or compact, there result some interesting properties of Nash equilibria. In particular, a game being Euclidean implies specific topological properties of the set of Nash equilibria of that game, while compactness and additional convexity conditions guarantee the existence of Nash equilibria of such a game. For a proof of the first insight I refer to Section 1.5.1.

Theorem 1.2

Let $\Gamma = \langle N, \mathbf{A}, \pi \rangle \in \mathbb{G}_E^N$ be a Euclidean game on the player set N . Then the set of Nash equilibria $NE(\Gamma) \subset \mathbf{A}$ is a closed set in the Euclidean topology on the profile space \mathbf{A} .
 If additionally Γ is a compact Euclidean game, then the set of Nash equilibria $NE(\Gamma) \subset \mathbf{A}$ is a compact set in the Euclidean topology on the profile space \mathbf{A} .

Existence of Nash equilibrium Nash (1951) seminally showed that Nash equilibria exist for a particular class of normal form games that can be denoted as mixed extensions of games with finite action sets. This was preceded by results published by Debreu (1952), Fan (1952), and Glicksberg (1952) that generalised Nash’s seminal result. We state this result as follows. For a proof of this result I refer to the last section of this chapter. For the more general existence result we recall some mathematical definitions.



Mathematical notes For the statement of Existence Theorem 1.3, we have to introduce some auxiliary mathematical concepts. For that purpose, let $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ be some Euclidean game such that for every player $i \in N$ the action set A_i is a nonempty closed subset of a finite dimensional Euclidean space and the payoff function π is continuous on \mathbf{A} . For this setting, we recall the following mathematical notions on the concavity of payoff functions.

- For $i \in N$ the payoff function π_i is **concave** in a_i if for all $a_i, b_i \in A_i$ and $a_{-i} \in \prod_{j \neq i} A_j$ it holds that $\pi_i(\lambda a_i + (1 - \lambda)b_i, a_{-i}) \geq \lambda \pi_i(a_i, a_{-i}) + (1 - \lambda)\pi_i(b_i, a_{-i})$.¹⁴
- Finally, for $i \in N$ the payoff function π_i is **quasi-concave** in a_i if for all $a_i, b_i \in A_i$ and $a_{-i} \in \prod_{j \neq i} A_j$ it holds that $\pi_i(\lambda a_i + (1 - \lambda)b_i, a_{-i}) \geq \min \{ \pi_i(a_i, a_{-i}), \pi_i(b_i, a_{-i}) \}$.

For the mathematical analysis of concave functions and related subjects, I refer to Rockafellar (1970). □

We can now state the Debreu-Fan-Glicksberg existence theorem. For a complete proof of this (historically) important result I refer to Section 1.5.2.

¹⁴I note that every concave function on a Euclidean space is continuous. This implies that all games consisting of a Euclidean pre-game endowed with a concave payoff function are actually Euclidean.

Theorem 1.3 (Existence of Nash equilibrium)

Let $\Gamma = \langle N, \mathbf{A}, \pi \rangle \in \mathbb{G}_C^N$ be a compact Euclidean game such that the following additional conditions are satisfied:

- (i) For every player $i \in N$ the action set A_i is a nonempty, compact and convex subset of a finite dimensional Euclidean space;
- (ii) For every player $i \in N$ the payoff function $\pi_i: \mathbf{A} \rightarrow \mathbb{R}$ is continuous in every profile $a \in \mathbf{A}$ and quasi-concave in every action $a_i \in A_i$.

Then the game Γ admits at least one Nash equilibrium.



The literature on strategic form games has developed more general Nash equilibrium existence results for non-Euclidean action sets. These contributions are usually mathematically much more advanced and rely on arguments regarding more general topological spaces. This textbook is not really the proper venue to provide an overview of these results.

1.3 Introducing potentials: Potential form games

The idea of a potential function was first applied in the 19th century in physics and mathematics. The main idea behind potential theory is that one can analyse certain mathematical constructs and models using a single function—referred to as the *potential function*. In 19th century physics potential theory was mainly used to summarise different forces in a physical setting in a single potential function. This mainly referred to the modelling of gravity and electrostatic forces through potential functions.

In game theory, potential functions have been used to describe payoff information in both non-cooperative normal form games and cooperative characteristic-function games. The two seminal contributions are Hart and Mas-Colell (1989), on cooperative games, and Monderer and Lloyd S Shapley (1996a), on normal form games. This section and the subsequent chapters pursue the non-cooperative line initiated by Monderer and Shapley; a fuller historical discussion is deferred to the notes to this chapter.

Non-cooperative games and potential functions The main idea is that a potential function summarises or captures essential information about a corresponding payoff structure on this pre-game. This is implemented through a potential function that assigns real numbered values to all action profiles in the underlying pre-game. Hence, the essential payoff structural features of a normal form game can be represented by such a potential function. More generally, we can introduce the notion of a potential function on any pre-game as set out in the next definition.

Definition 1.7 (Potential form games)

Let $(A_i)_{i \in N}$ be some pre-game on the finite player set N . A function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ is referred to as a **potential function** on the pre-game $(A_i)_{i \in N}$.

The triple $\Gamma_\Psi = \langle N, \mathbf{A}, \Psi \rangle$, where $(A_i)_{i \in N}$ is a pre-game on player set N and $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ is a potential function on this pre-game, is denoted as a **potential form game**.



A potential form game can be interpreted as a standard normal form game in which all players have been assigned identical payoff functions, namely the potential function itself. Hence, $\Gamma_\Psi = \langle N, \mathbf{A}, \Psi \rangle$ can be interpreted as

a normal form game $\langle N, \mathbf{A}, \pi \rangle$ where $\pi_i = \Psi$ for all players $i \in N$. These particular type of games are also known as “coordination games” or “team games”.

As before we can introduce specific classes of potential form games. The most obvious class is the one that is based on the properties of Euclidean vector spaces.

Definition 1.8 (Euclidean potential form games)

Let $\Gamma_\Psi = \langle N, \mathbf{A}, \Psi \rangle$ be a potential form game with potential function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$.

- (i) The potential form game Γ_Ψ is **Euclidean** if the underlying pre-game $(A_i)_{i \in N}$ is Euclidean and the potential function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ is continuous on \mathbf{A} .
- (ii) The potential form game Γ_Ψ is **compact** if Γ_Ψ is Euclidean and the profile space $\mathbf{A} = \prod_{i \in N} A_i$ is a compact set in some finite dimensional Euclidean space.



A consequence of Proposition 1.1 is that every potential form game in which all players have finite action sets is Euclidean as well as compact. This is stated in the next corollary.

Corollary 1.1

Every finite potential form game is compact.



When a potential form game is infinite, it may be Euclidean or non-Euclidean. For a potential form game to be non-Euclidean, it is clear that actions need to be of rather complex a nature. This has already been discussed for Euclidean (pre-) games. The next example considers a nice quadratic potential form game.

Example 1.6 Consider an interactive decision situation with two players, represented by $N = \{1, 2\}$. Furthermore, take any pair of numbers $\alpha, M \in \mathbb{R}$ with $M > 0$. We now construct a potential form game $\Gamma_{(\alpha, M)}$ with a potential function Ψ_α .

In this potential form game, both players are assigned exactly the same action set given by $A_1 = A_2 = [-M, M]$. Hence, the profile space is given as $\mathbf{A} = [-M, M]^2 \subset \mathbb{R}^2$.

Second, we introduce a quadratic potential function $\Psi_\alpha: [-M, M]^2 \rightarrow \mathbb{R}$ by

$$\Psi_\alpha(a_1, a_2) = 2\alpha a_1 a_2 - a_1^2 - a_2^2$$

Note that the generated potential form game $\Gamma_{(\alpha, M)} = \langle N, [-M, M]^2, \Psi_\alpha \rangle$ is compact. In fact, the potential function Ψ_α is clearly twice differentiable on the relative interior of the profile space $\mathbf{A}^\circ = (-M, M)^2$. ♦

1.3.1 Defining equilibrium in potential form games

The concept of Nash equilibrium can be applied to the class of potential form games, interpreted as normal form games with all players having identical payoff functions. Due to the particular payoff structure of a potential form game—represented by a single potential function—these Nash equilibria are exactly the class of “saddle points” of the potential function.


We simply refer to these saddle points as *equilibria* of the corresponding potential form game.¹⁵ The notion of equilibrium for potential form games is formalised in the next definition.

¹⁵Indeed, in a potential form game, there is no reason to denote an equilibrium point as a “Nash equilibrium” to distinguish it from other forms of equilibrium. As shown later in this text, these equilibrium points are representing other equilibrium conceptions as well.

Definition 1.9

Let $\Gamma_\Psi = \langle N, \mathbf{A}, \Psi \rangle$ be a potential form game with potential function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$.
 A profile $\hat{a} \in \mathbf{A}$ is an **equilibrium** in Γ_Ψ if for every player $i \in N$:

$$\Psi(b_i, \hat{a}_{-i}) \leq \Psi(\hat{a}) \quad \text{for every alternative action } b_i \in A_i. \quad (1.12)$$

The set of all equilibria in Γ_Ψ is denoted by $E(\Gamma_\Psi) \subset \mathbf{A}$. 

We note that from the definition there emerge three types of equilibria of Euclidean potential form games. First, there are the pure saddle points of the potential function. These are stationary profiles¹⁶ that are neither a minimum nor a maximum of the potential function. Second, all maxima of the potential functions are equilibria. Third, there might be points on the boundary of the profile space as a subset of a Euclidean vector space that exhibit the properties of a saddle point.

The first two types of equilibria are quite common in Euclidean potential form games. In fact, the set of maxima of the potential function is considered to be a “refinement” of the equilibrium concept for potential form games. This is subject to the next notes.



Mathematical notes Consider a Euclidean potential form game $\Gamma_\Psi = \langle N, \mathbf{A}, \Psi \rangle$ such that the potential function Ψ is twice differentiable on the relative interior of \mathbf{A} , which we denote by \mathbf{A}° .¹⁷ For every profile $a \in \mathbf{A}^\circ$ these first and second order derivatives are denoted by

$$D\Psi(a) = \left(\frac{\partial}{\partial a_1} \Psi(a), \dots, \frac{\partial}{\partial a_n} \Psi(a) \right)$$

and

$$D^2\Psi(a) = \begin{bmatrix} \frac{\partial^2}{\partial a_1^2} \Psi(a) & \cdots & \frac{\partial^2}{\partial a_1 \partial a_n} \Psi(a) \\ \vdots & \ddots & \vdots \\ \frac{\partial^2}{\partial a_n \partial a_1} \Psi(a) & \cdots & \frac{\partial^2}{\partial a_n^2} \Psi(a) \end{bmatrix},$$

respectively. The matrix $D^2\Psi(a)$ of partial second derivatives is also referred to as the **Hessian** of the potential function Ψ in $a \in \mathbf{A}$. In case that Ψ is twice continuously differentiable, the Hessian is a symmetric $n \times n$ matrix.

Using these concepts, we can now characterise some sufficient conditions for some profile in the relative interior of the profile space to be an equilibrium or a potential maximiser.

- A **stationary profile** of the potential function Ψ is any profile $a \in \mathbf{A}^\circ$ at which the derivative is zero, i.e., such that $D\Psi(a) = 0$. Stationarity of a profile is a prerequisite for a profile to be a saddle point or a maximum or a minimum of the potential function.¹⁸ As identified, equilibria are either saddle points or maximisers of the potential function Ψ . Clearly, minimisers have to be weeded out from the stationary points of Ψ .

¹⁶A profile $a \in \mathbf{A}$ is *stationary* if the derivative of the potential function is zero in that profile, i.e., $D\Psi(a) = 0$.

¹⁷Here we use the notation that S° refers to the relative interior of a set $S \subset \mathbb{R}^k$ in the k -dimensional Euclidean real vector space.

¹⁸For details on the calculus of functional optimisation I also refer to Apostol (1974), Chapter 13, and the mathematical appendix of Jehle and Reny (2011), A2.2.

- *Secondary conditions can help us find equilibria in the potential form game. Indeed, a profile $\hat{a} \in \mathbf{A}^\circ$ is an **equilibrium** of Γ_Ψ if \hat{a} is a stationary point of the potential function Ψ such that for every player $i \in N$ the equilibrium action $\hat{a}_i \in A_i^\circ$ is a global maximiser for the i -th dimensional restriction of Ψ given by $\Psi(\hat{a}_1, \dots, \hat{a}_{i-1}, \cdot, \hat{a}_{i+1}, \dots, \hat{a}_n): A_i \rightarrow \mathbb{R}$.*

This is the case if for every $i \in N$:

$$\frac{\partial^2}{\partial a_i^2} \Psi(\hat{a}) < 0.$$

The secondary conditions on the second order derivatives $\frac{\partial^2}{\partial a_i^2} \Psi(\hat{a})$ for $i \in N$ guarantee that all one-dimensional restrictions of Ψ at \hat{a} are locally concave one-dimensional functions, implying that any stationary point on that one-dimensional restriction is actually a local maximiser.

- *Finally, using additional secondary conditions, a profile $\hat{a} \in \mathbf{A}^\circ$ can also be identified as a (local) **maximiser** of Ψ if \hat{a} is a stationary point of Ψ such that the second order derivative or Hessian $D^2 \Psi(\hat{a})$ is negative definite for the stationary point profile $\hat{a} \in \mathbf{A}^\circ$.¹⁹ The mathematical test for $D^2 \Psi(\hat{a})$ being negative definite is that its determinant $\det D^2 \Psi(\hat{a}) > 0$ is positive and $\frac{\partial^2 \Psi}{\partial a_i^2}(\hat{a}) < 0$ for every $i \in N$. The latter condition on the Hessian implies that the potential function Ψ is locally concave in a neighbourhood of \hat{a} on \mathbf{A} , implying that any interior stationary point is necessarily a local maximiser. Note that concavity of the potential function refers here to a multi-dimensional property rather than the one-dimensional concavity referred to in the above notes about equilibria. From the above, it should be clear that every (local) maximiser of Ψ is definitely an equilibrium of Γ_Ψ .*

*From the above we conclude that it might be that certain equilibrium profiles of Ψ are not (global) maximisers. Therefore, the set of maximisers of Ψ forms a subset of $E(\Gamma_\Psi)$, possibly even a strict subset. Hence, potential function maximisers form an equilibrium **refinement** of the standard equilibrium concept for potential form games. \square*

These notes only concern the profiles in the relative interior \mathbf{A}° of the profile space \mathbf{A} . There might be equilibria and maximisers of the potential function Ψ on the boundary of the profile space $\partial \mathbf{A}$. The next example exhaustively analyses the quadratic potential function in a two-player potential form game already introduced in Example 1.6.

Example 1.7 Consider the potential form game $\Gamma_{(\alpha, M)} = \langle N, [-M, M]^2, \Psi_\alpha \rangle$ introduced in Example 1.6, where the profile space is given by $\mathbf{A} = [-M, M]^2 \subset \mathbb{R}^2$ and the potential function is the quadratic form $\Psi_\alpha(a_1, a_2) = 2\alpha a_1 a_2 - a_1^2 - a_2^2$.

First, we note that on $\mathbf{A}^\circ = (-M, M)^2$ it holds that

$$D \Psi_\alpha(a) = (2\alpha a_2 - 2a_1, 2\alpha a_1 - 2a_2)$$

and

¹⁹ A $k \times k$ -matrix T is said to be *negative semi-definite* if all its eigenvalues are non-positive and *negative definite* if all its eigenvalues are strictly negative. Hence, for all non-zero vectors $x \in \mathbb{R}^k \setminus \{0\}$ and real values $\lambda \in \mathbb{R}$ with $Tx = \lambda x$ it holds that $\lambda < 0$. Equivalently, the matrix T is negative definite if and only if the corresponding quadratic form $Q_T(x) = x^T T x < 0$ is negative for every $x \in \mathbb{R}^k$.

$$D^2 \Psi_\alpha(a) = \begin{bmatrix} -2 & 2\alpha \\ 2\alpha & -2 \end{bmatrix}.$$

On the relative interior $\mathbf{A}^\circ = (-M, M)^2 \subset \mathbb{R}^2$ we can now compute all stationary points of the potential function Ψ_α for any $\alpha \in \mathbb{R}$. Indeed, all stationary points are determined by the system of two linear equations

$$\begin{cases} 2\alpha a_2 - 2a_1 = 0 \\ 2\alpha a_1 - 2a_2 = 0 \end{cases}$$

This is equivalent to

$$\begin{cases} \alpha a_2 = a_1 \\ \alpha a_1 = a_2 \end{cases}$$

This leads to the following insights regarding the nature of the stationary points of Ψ_α :

$$\begin{array}{ll} \hat{a} = (0, 0) & \text{is a stationary point for all } \alpha \in \mathbb{R} \\ \hat{a}_\beta = (\beta, \beta) \text{ with } -M < \beta < M & \text{are stationary points for } \alpha = 1 \\ \hat{a}_\gamma = (\gamma, -\gamma) \text{ with } -M < \gamma < M & \text{are stationary points for } \alpha = -1 \end{array}$$

From the second order derivatives we note immediately that all stationary points in $(-M, M)^2$ are actually equilibria, since $\frac{\partial^2}{\partial a_1^2} \Psi_\alpha = \frac{\partial^2}{\partial a_2^2} \Psi_\alpha = -2 < 0$ on $(-M, M)$. Hence, the partial second order derivatives being negatives implies that all stationary points are maximisers of the two one-dimensional restrictions of the potential function Ψ_α for any $\alpha \in \mathbb{R}$.

Next, we note that the second order derivative of Ψ_α is independent of the profile at which it is evaluated. Therefore, the determinant of the second order derivative of Ψ_α is computed as $\det D^2 \Psi_\alpha(a) = (-2)^2 - (2\alpha)^2 = 4 - 4\alpha^2$ for any $a \in (-M, M)^2$. Hence, with the previous, a stationary point $\hat{a} \in (-M, M)^2$ is a maximiser of Ψ_α if and only if $D^2 \Psi_\alpha(a) \leq 0$ for all $a \in (-M, M)^2$ if and only if $\det D^2 \Psi_\alpha(a) \geq 0$ for all $a \in (-M, M)^2$ if and only if $\alpha^2 \leq 1$ if and only if $-1 \leq \alpha \leq 1$.

We conclude that Ψ_α has a unique (strict) global maximiser $\hat{a} = (0, 0)$ for $-1 < \alpha < 1$. This is also the unique interior equilibrium of $\Gamma_{(M, \alpha)}$ for $M > 0$ and $-1 < \alpha < 1$.

On the other hand, the potential form game Ψ_α also admits a number of equilibria on the boundary of the profile space $[-M, M]^2$. In particular, we identify the following additional equilibria on this boundary:

- $\alpha < -1$: Besides the equilibrium $\hat{a} = (0, 0)$, which is a true saddle point of Ψ_α , we identify $(-M, M)$ and $(M, -M)$ as global maximisers of Ψ_α with $\Psi_\alpha(M, -M) = \Psi_\alpha(-M, M) = -2(\alpha + 1)M^2 > 0$ and, therefore, as additional equilibria.

We emphasise here that in these cases $\hat{a} = (0, 0)$ is an equilibrium that is *not* a (global) maximiser of the potential function Ψ_α .

- $\alpha = -1$: All profiles $(\gamma, -\gamma)$ with $-M \leq \gamma \leq M$ are identified as equilibria of $\Gamma_{(M, \alpha)}$, where all $(\gamma, -\gamma)$ are global maximisers of Ψ_α with $\Psi_\alpha(\gamma, -\gamma) = 0$.
- $-1 < \alpha < 1$: The profile $\hat{a} = (0, 0)$ is the unique equilibrium of $\Gamma_{(M, \alpha)}$ as well as the unique maximiser of Ψ_α with $\Psi_\alpha(0, 0) = 0$.

- $\alpha = 1$: All profiles (β, β) with $-M \leq \beta \leq M$ are identified as equilibria of $\Gamma_{(M, \alpha)}$, where all (β, β) are global maximisers of Ψ_α with $\Psi_\alpha(\beta, \beta) = 0$.
- $\alpha > 1$: Besides the equilibrium $\hat{a} = (0, 0)$, which is a true saddle point of Ψ_α , we identify $(-M, -M)$ and (M, M) as global maximisers of Ψ_α , with $\Psi_\alpha(M, M) = \Psi_\alpha(-M, -M) = 2(\alpha - 1)M^2 > 0$ and, therefore, as additional equilibria.

Again we note here that in these cases $\hat{a} = (0, 0)$ is an equilibrium that is *not* a (global) maximiser of the potential function Ψ_α .

This simple example of a potential form game shows that a wide variety of equilibria can emerge. In particular, we note for some values of the parameter α all equilibria are potential maximisers, while for other parameter values of α the set of potential maximisers is a strict subset of the set of equilibria $E(\Gamma_{(M, \alpha)})$. In the latter case, the subset of potential function maximisers forms a proper refinement of the set of the potential form game equilibria. \blacklozenge

1.3.2 Nash potential games

The fundamental idea behind potential functions and potential form games is that relevant information about a normal form game can be “summarised” in a corresponding potential function, so that the associated potential form game represents the original.

The benefit of such a construction is that relevant equilibria can be computed much more easily in the potential form game than in the original normal form game. This is particularly useful in dynamic social decision situations that require continuous updating of equilibrium actions.

In this and the next chapters we explore various constructs that connect normal form games with corresponding potential form games, that is, classes of normal form games that can, to varying degrees, be represented through a well-constructed potential function.

The most basic such construct is the representation of equilibria only: a normal form game is a “Nash potential game” if some potential function generates a set of equilibria equal to the set of Nash equilibria of the original game. This is formalised in the next definition, which originates in Voorneveld (1999, Section 7.5).

Definition 1.10 (Nash potential games)

A normal form game $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ is a **Nash potential game** if there exists some potential function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ such that for the corresponding potential form game $\Gamma_\Psi = \langle N, \mathbf{A}, \Psi \rangle$ it holds that $E(\Gamma_\Psi) = NE(\Gamma)$, i.e., $\hat{a} \in \mathbf{A}$ is a Nash equilibrium of Γ if and only if \hat{a} is an equilibrium of Γ_Ψ . \clubsuit

It is appropriate to explore a few examples of potential form games and to investigate the form that these constructed potential functions assume. This is the subject of the next example, which discusses some of the examples considered in Examples 1.1 and 1.3.

Example 1.8 We again consider some of the games explored in Examples 1.1, 1.3 and 1.4. In particular, we investigate whether these games are Nash potential games and what form the corresponding potential functions take. For that purpose let the player set be given by $N = \{1, 2\}$ as before.

- Recall the **Rock-Paper-Scissors game** discussed before. Both players had identical action sets $A_1 = A_2 = \{R, P, S\}$ with payoffs given in the matrix representation

	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

This finite game has no Nash equilibrium.

For this RPS-game to be a Nash potential game, one needs to construct a potential function $\Psi: A_1 \times A_2 \rightarrow \mathbb{R}$ such that $E(\Psi) = \emptyset$. I argue that this is impossible.

Indeed, any assignment of numerical values Ψ to fields in the matrix depicted above will assign a highest or “maximal” value to at least one of these fields. Hence, Ψ would attain a maximum. So, we conclude that the corresponding maximisers are equilibria of the potential form game based on this potential function Ψ .

In fact, one can now conclude that, since *every* potential function on *any* finite profile set always attains a maximum, the corresponding maximisers are equilibria of the corresponding potential form game. Hence, the set of equilibria of *every* finite potential form game is always non-empty.

- (ii) Consider the simple voting game discussed in Examples 1.1, 1.3 and 1.4(iii). The variations of this game—depending on the exact voting mechanism considered—are example of *finite normal form games that admit Nash equilibria*. For these games, one can formulate a potential function showing that these games are Nash potential games. Indeed, for such a finite game Γ with $NE(\Gamma) \neq \emptyset$ consider the potential as the indicator or characteristic function of the set of Nash equilibria:

$$\Psi(a) = \chi_{NE(\Gamma)}(a) = \begin{cases} 1 & \text{if } a \in NE(\Gamma) \\ 0 & \text{if } a \in \mathbf{A} \setminus NE(\Gamma) \end{cases} .$$

It is clear that all Nash equilibrium profiles $\hat{a} \in NE(\Gamma) \neq \emptyset$ are the potential maximisers and all other profiles $a \notin NE(\Gamma)$ are potential minimisers. This implies that $NE(\Gamma) = E(\Gamma_\Psi)$, showing that every finite game that admits Nash equilibria is actually a Nash potential game.

- (iii) Next, we consider the collective good provision game. Recall that we have n players, action sets $A_i = [0, M_i]$ with $M_i \geq 0$ for every $i \in N$, and payoffs given by

$$\pi_i(a) = \begin{cases} \alpha_i - a_i & \text{if } \sum_N a_i \geq C \\ \beta_i & \text{if } \sum_N a_i < C \end{cases}$$

where $\alpha_i, \beta_i \geq 0$ and $C \geq 0$ the total cost of providing the collective good or project. The actions are interpreted as voluntary contributions to the provision of the collective project.

We can now fully characterise all Nash equilibria in this game. For that, we introduce the *maximal contribution to the collective project* that a player is willing to make. For $i \in N$ this is given by

$$\bar{a}_i = \min \{M_i, \max \{0, \alpha_i - \beta_i\}\} \in A_i = [0, M_i].$$

In particular, $\bar{a}_i = 0$ if $\beta_i \geq \alpha_i$, which corresponds to the player having no benefit from the collective

project. Using this maximal contribution we can characterise the Nash equilibria:

- If $\sum_N \bar{a}_i \geq C$, then a profile $\hat{a} \in \mathbf{A}$ is an equilibrium if and only if $0 \leq \hat{a}_i \leq \bar{a}_i$ and $\sum_N \hat{a}_i = C$.
- If $\sum_N \bar{a}_i < C$, then a profile $\hat{a} \in \mathbf{A}$ is an equilibrium if and only if $\sum_N \hat{a}_i < C$.

We derive from this the set-theoretic description of the equilibria of this provision game. Define

$$\mathbf{A}' = \prod_{i \in N} [0, \bar{a}_i] \subseteq \mathbf{A} \quad \text{and} \quad \mathbf{A}'_C = \left\{ a \in \mathbf{A}' \mid \sum_{i \in N} a_i \geq C \right\}$$

Clearly, \mathbf{A}' is the set of profiles that are feasible in terms of the players' willingness to contribute to the collective project and \mathbf{A}'_C is the set of the feasible profiles that actually lead to the provision of the collective project. Now the Nash equilibria are exactly the feasible profiles in which the collective project is provided through *exact* financing:

$$\text{NE} = \left\{ a \in \mathbf{A}' \mid \sum_{i \in N} a_i = C \right\}$$

which forms the upper boundary of \mathbf{A}'_C . Clearly, $\text{NE} \neq \emptyset$ if and only if $\sum_N \bar{a}_i \geq C$.

- (iv) Turning to the Cournot duopoly, I use the formulation from Example 1.4(i). There the payoff structure was generalised to be given by

$$\begin{aligned} \pi_1(Q_1, Q_2) &= \alpha Q_1 - \beta Q_1(Q_1 + Q_2) - c_1(Q_1) \\ \pi_2(Q_1, Q_2) &= \alpha Q_2 - \beta Q_2(Q_1 + Q_2) - c_2(Q_2) \end{aligned}$$

where c_1 and c_2 are convex cost functions for both firms in the market.

This game has a structure that allows a much more advanced summary of the payoff structure into a single potential function. Indeed, we can devise a potential function $P: \mathbb{R}_+^2 \rightarrow \mathbb{R}$ that *exactly* represents the payoff differences if both firms change their respective output level. Such an *exact potential function* can be constructed as

$$P(Q_1, Q_2) = \alpha(Q_1 + Q_2) - \beta Q_1 Q_2 - \beta Q_1^2 - \beta Q_2^2 - c_1(Q_1) - c_2(Q_2).$$

Now we note that²⁰

$$P(Q_1, Q_2) - P(Q'_1, Q_2) = \pi_1(Q_1, Q_2) - \pi_1(Q'_1, Q_2) \quad \text{for all } Q_1, Q'_1 \geq 0 \text{ and } Q_2 \geq 0$$

as well as

$$P(Q_1, Q_2) - P(Q_1, Q'_2) = \pi_2(Q_1, Q_2) - \pi_2(Q_1, Q'_2) \quad \text{for all } Q_2, Q'_2 \geq 0 \text{ and } Q_1 \geq 0.$$

²⁰Showing these equalities is left to the interested reader.

This implies immediately that the equilibria for the corresponding potential form game are exactly the same as the Nash equilibria of this duopoly. Hence, the duopoly is a Nash potential game for the exact potential function P .

On the other hand, the potential function P captures much more information of this game than just the Nash equilibria. Indeed, the function P captures the complete payoff structure. This can be exploited to use algorithms to easily find the Nash equilibria of the duopoly through the potential function. This is the subject of the next chapter in this text.

These examples show clearly the intricacies of constructing potential functions with the pre-game structures of these games to represent the payoffs introduced. ◆

The class of potential form games is very broad as confirmed by the next theorem. This theorem confirms that essentially only the class of finite games that do *not* admit any Nash equilibria is not a subclass of the broad class of Nash potential games. The next theorem generalises Theorem 7.9 of Voorneveld (1999).

Theorem 1.4

Let $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ be some normal form game. Then the following statements hold:

- (a) If Γ admits at least one Nash equilibrium, i.e., $\text{NE}(\Gamma) \neq \emptyset$, then Γ is a Nash potential game.
- (b) If there is some player $j \in N$ with an infinite action set A_j and Γ admits no Nash equilibrium, i.e., $\text{NE}(\Gamma) = \emptyset$, then Γ is a Nash potential game.
- (c) If for every player $i \in N$ the action set A_i is a compact set and Γ is a Nash potential game for some continuous potential function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$, then Γ admits at least one Nash equilibrium, $\text{NE}(\Gamma) \neq \emptyset$.



A proof of Theorem 1.4 is developed in Section 1.5.3 below.

The first assertion of Theorem 7.9 in Voorneveld (1999) now follows immediately from Theorem 1.4(a) and (c), since all finite sets are trivially compact and all functions on finite sets are trivially continuous. This is stated as the next corollary.

Corollary 1.2 (Theorem 7.9, Voorneveld, 1999)

If $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ is a finite normal form game, then Γ is Nash potential game if and only if Γ admits at least one Nash equilibrium, $\text{NE}(\Gamma) \neq \emptyset$.



Properties of Euclidean Nash potential games are generally stronger than arbitrary Nash potential games. The next theorem addresses the existence of continuous potential functions for Euclidean games that admit Nash equilibria. A proof of Theorem 1.5 is developed in Section 1.5.4.

Theorem 1.5

Let $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ be a Euclidean normal form game that admits Nash equilibria, i.e., $\text{NE}(\Gamma) \neq \emptyset$. Then there exists a continuous potential function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ such that $E(\Gamma_\Psi) = \text{NE}(\Gamma)$.



Some remarks on potentials and Nash potentials I have introduced the notion of a potential game as a vehicle for further development in the main body of this text. The notion also supports the study of non-cooperative

normal form games that satisfy much stronger properties, for which the relationship between the payoff structure of the game and its potential-function representation can be sharpened considerably.

The class of Nash potential games is the broadest class that can be related to a representative potential form game: for every such game, a potential function links its Nash equilibria to the equilibria of the corresponding potential form representation.

This is a rather weak property, as witnessed by the breadth of the class established in Theorem 1.4. The informational benefits of the representation are correspondingly limited: the construction in the next section shows that the potential function carries no information about the payoff structure of the represented game beyond the topological or metric properties of its Nash equilibrium set.

The next chapters develop more insightful representations of normal form games through potential functions. In the next chapter I set out the development of so-called exact potential games in which there is maximal information about the payoff structure represented through the corresponding potential function. Clearly, the class of exact potential games is rather small, signifying that the requirements are quite demanding and that the informational content of the potential function representation of the game is substantial. This is also clear from the number of strong properties that these games satisfy. This ranges from the structure of Nash equilibria through the possibility to compute Nash equilibria through simple learning algorithms. These insights are explored throughout this text.

Subsequent chapters introduce further classes of normal form games with less informative representations through potential functions. This results in larger classes of games, indicating the less informative value of these potential function representations. The properties of these classes of games are weaker as well, leading to less straightforward analysis of the structure of their Nash equilibria and increased complexity of learning algorithms to compute equilibria.

1.4 Notes

I close this chapter by stepping back from the formal apparatus of Sections 1.2 through 1.3 to reflect on three questions that the preceding material has raised but left unanswered. First, why do we call the function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ a *potential*? Second, how did the concept of a potential enter game theory, and by what route did it arrive at the form in which Section 1.3 now presents it? Third, for what classes of problems has the concept proved decisive? The discussion that follows is necessarily selective. I point forward to the chapters in which each strand is developed in detail.

How “potential” got its name The term “potential” is inherited from classical mathematical physics, where it designates a scalar field whose gradient encodes a vector field of forces. The Preface to this book reproduces the four meanings that Samuel Johnson recorded in his *Dictionary of the English Language* of 1755 (Johnson 1755). Across those senses a common thread is discernible: a potential describes what *could* happen in a system—its latent capacity for change—rather than what has actually been realised. It is precisely this reading that the mathematical usage inherits.

The technical notion of a potential function is due to the programme of eighteenth- and nineteenth-century celestial mechanics and electrostatics. Lagrange’s work on the three-body problem introduced, as early as 1773, a scalar function whose gradient gives the gravitational force. Laplace subsequently formulated the equation $\nabla^2 V = 0$ for the gravitational potential in empty space. The term itself—“potential function”—is commonly attributed to George Green’s *Essay on the Application of Mathematical Analysis to the Theories of Electricity*

and *Magnetism* (Green 1828), and was subsequently systematised by Gauss and codified in the twentieth-century treatise of Kellogg (1929). The common structural feature, across all of these settings, is that a vector field of forces is represented as the gradient of a single scalar field, so that differences in that scalar between two points measure the work done against the field in moving between them.

This structural analogy licenses the use of the term in the present setting. In a potential form game $\Gamma_\Psi = \langle N, \mathbf{A}, \Psi \rangle$, the “force” that acts on player i at a profile $a \in \mathbf{A}$ is the collection of payoff differences $\pi_i(b_i, a_{-i}) - \pi_i(a_i, a_{-i})$ generated by her unilateral deviations. When the game admits an exact potential, the scalar field Ψ stands to those payoff-difference “forces” in the same relation that the electrostatic potential stands to the electric force field: its differences along unilateral-deviation axes reproduce the payoff differences exactly. The analogy is more than cosmetic. It is responsible for the conservation-law character of exact potential games—the absence of payoff-improving cycles—and it underlies the Helmholtz-style orthogonal decomposition of finite games into potential, harmonic, and non-strategic components that Chapter 2 will develop.

1.4.1 A short history of potentials in game theory

The systematic use of potential-function methods in game theory emerged comparatively late. I distinguish three phases.

Phase I: prehistory (1956–1988) An implicit use of a potential function appears already in the transportation-economics monograph of Beckmann, McGuire, and Winsten (1956). These authors show that the user-optimal traffic flow, now known as the Wardrop equilibrium, can be characterised as a minimiser of a convex integral functional on the space of route flows. The functional in question is a potential in everything but name, though at the time it was not recognised as such.

The first formal instance of an exact potential in a non-cooperative game is due to Rosenthal (1973). For the class of finite congestion games he constructed an explicit real-valued function on action profiles whose decrease along any unilateral improvement step is precisely the improvement of the deviating player. Rosenthal called it simply “a function”; the recognition that this function is the prototype of a general notion, and that it deserves a name, came only two decades later.

Phase II: foundations (1989–1996) The foundations of the subject rest on two seminal papers.

Hart and Mas-Colell (1989), writing on cooperative game theory, introduced a potential P defined on the family of subgames of a transferable-utility cooperative game, and proved that the Shapley value of any subgame equals the vector of discrete marginal derivatives of P . This is a deep result. It establishes that the axiomatically characterised Shapley value of Lloyd S Shapley (1953) is itself the gradient of a scalar field—the first instance of a game-theoretic payoff rule being derived from a potential. The paper is the locus classicus for the introduction of the term “potential” into game theory in an explicit and systematic form.

Monderer and Lloyd S Shapley (1996a), building on the cooperative-game terminology of Hart and Mas-Colell, transposed the idea to non-cooperative games. They defined exact, weighted, ordinal, and generalised-ordinal potential functions for finite normal form games; proved the equivalence between exact potentiality and the absence of improvement cycles; showed that the finite improvement property characterises the generalised-ordinal class; and established that the class of finite congestion games of Rosenthal (1973) coincides exactly with the class of finite exact potential games. This paper is the foundational reference for the non-cooperative theory that the remainder of this book develops.

Phase III: maturation and branching (1996–present) The non-cooperative strand opened by Monderer and

Shapley has since grown into several distinct research programmes. A selective sketch must suffice here. Neyman (1997), Ui (2001), and more recently Cao, Tan, and Zhou (2024) and Swenson, Murray, and Kar (2020) have clarified the fine structure of the equilibrium set of exact potential games, including the role of the potential maximiser as a refinement and selection device. This material is developed in Chapters 2 and 3.

The algebraic and geometric structure of *finite* potential games was reorganised by Candogan, Menache, et al. (2011) through a Helmholtz-style orthogonal decomposition of the finite-game vector space, and independently by Cheng (2014) through the semi-tensor-product formalism. Both threads are taken up in Chapter 4.

The dynamic strand—fictitious play, best-response and better-response dynamics, logit and no-regret learning—exploits the potential as a Lyapunov function, a theme that runs through the dynamics chapter. Finally, the cooperative and non-cooperative threads were reconnected by Ui (2000), who showed that a non-cooperative game admits a Monderer-Shapley exact potential if and only if the Shapley value of an associated cooperative game coincides with the gradient of that potential (Section 2.2.2).

1.4.2 Main fields of application

The payoff of a potential-function representation, in every setting in which it is available, is that a single scalar field simultaneously delivers two of the central questions of non-cooperative game theory: existence of pure-strategy Nash equilibrium (a maximiser of the potential, if one exists, is an equilibrium) and convergence of a large class of learning dynamics (the potential serves as a Lyapunov function). The following is a brief selection of application areas in which one or both of these payoffs has proved decisive. Subsequent chapters treat each of them in detail.

Congestion and routing games The paradigmatic application, and the motivating example of Rosenthal (1973).

Users share scarce resources—road segments, communication links, server capacity—whose cost to each user is increasing in the aggregate number of users. Rosenthal’s potential function establishes the existence of a pure-strategy Nash equilibrium in every finite congestion game, and it is the technical engine behind both the price-of-anarchy bounds of Roughgarden (2006) and the PLS-completeness of pure-equilibrium computation proved by Fabrikant, Papadimitriou, and Talwar (2004). This application runs as a red thread through the later applied chapters of this book.

Oligopoly and aggregative games Cournot competition with linear or quadratic inverse demand and constant or quadratic costs is an exact potential game, as observed seminally by Slade (1994) and Monderer and Lloyd S Shapley (1996a). The potential is a weighted sum of individual profits adjusted by a cross-term in aggregate output. The existence and uniqueness theorems of aggregative-games theory—Folmer and Mouche (2004) and, recently, Nocke and Schutz (2025)—are at root potential-function arguments. A Cournot running example is used throughout Chapters 2 and 3 to make the abstract theory concrete.

Learning, dynamics and multi-agent control In any ordinal potential game, best-response and better-response dynamics terminate at a Nash equilibrium; in exact potential games, fictitious play and a wide range of no-regret learning algorithms converge (Anagnostides et al. 2025; Monderer and Lloyd S. Shapley 1996b). This convergence property has made potential games the natural testbed for learning dynamics in multi-agent systems. It has also motivated an inverse programme in control engineering: given a desired global objective, design player-level utility functions such that the resulting game has the objective as its potential (Marden 2012; Marden, Arslan, and Shamma 2009; Marden and Shamma 2015). The dynamics chapter

takes up both directions.

Decision-theoretic foundations A more surprising connection was established recently by Ismail and Peeters (2024). These authors show that a finite two-player symmetric game is a potential game if and only if the preference relation that its payoff function represents, on lotteries over the common action set, satisfies the von Neumann-Morgenstern expected-utility axioms. This places the class of finite symmetric potential games in exact correspondence with the decision-theoretic foundations of expected-utility theory, and provides a principled justification—distinct from equilibrium selection—for singling out the potential maximiser. The result is developed in Chapter 4.

Machine learning and generative models A second, and at first sight unexpected, application has emerged from the recent machine-learning literature. Policy-gradient methods in multi-agent reinforcement learning, and certain training schemes for generative adversarial networks, exhibit exact or approximate potential structure in identifiable special cases; where the structure is present, the attendant convergence guarantees follow from the potential-function arguments of classical game theory (Roughgarden 2016). The near-potential framework of Candogan, Ozdaglar, and Parrilo (2013) is the natural technical bridge between the clean exact-potential case and the approximate-potential games that arise in these machine-learning settings. A brief treatment appears in the computational chapter.

A closing remark The pattern running through each of these applications is the same. A scalar potential function, where it exists, collapses the analysis of strategic interaction onto the analysis of a single optimisation problem, and it does so without discarding the underlying strategic structure. It is this combination—strategic fidelity together with analytical tractability—that explains the continued relevance of potential-function methods across disciplines as distant as transportation economics, control engineering, and machine learning. The next chapter begins the systematic development of the principal subclass within which this programme can be pursued: the exact potential games of Monderer and Shapley.

1.5 Proofs of Theorems and Propositions

1.5.1 Proof of Theorem 1.2

Let $\Gamma = \langle \mathbf{A}, \pi \rangle \in \mathbb{G}_E^N$ be a Euclidean game on the player set N and let $\text{NE}(\Gamma) \subset \mathbf{A}$ be the set of Nash equilibria of Γ . If $\text{NE}(\Gamma) = \emptyset$, it trivially follows that this set is compact and, therefore, closed.

Next, assume that $\text{NE}(\Gamma) \neq \emptyset$. To show that $\text{NE}(\Gamma)$ is a closed set, we show that every convergent sequence in $\text{NE}(\Gamma)$ has its limit contained in $\text{NE}(\Gamma)$.

Take some convergent sequence $(a^n)_{n \in \mathbb{N}}$ in $\text{NE}(\Gamma)$ with $a^n \rightarrow a \in \mathbf{A}$. Then by the fact that a^n is a Nash equilibrium of Γ , for every player $i \in N$ and every alternative action $b_i \in A_i$: $\pi_i(a^n) \geq \pi_i(b_i, a_{-i}^n)$. By taking the limit of (a^n) on both sides of this weak inequality, the continuity of player i 's payoff function π_i implies that $\pi_i(a) \geq \pi_i(b_i, a_{-i})$. Hence, since $i \in N$ and $b_i \in A_i$ are arbitrary, this shows that $a \in \text{NE}(\Gamma)$. Hence, $\text{NE}(\Gamma)$ is closed.

Finally, if the game Γ is compact, the compactness of \mathbf{A} implies that the closed set of Nash equilibria $\text{NE}(\Gamma)$ is actually compact. This completes the proof of Theorem 1.2.

1.5.2 Proof of Theorem 1.3

In this subsection we set out the standard proof of the existence of Nash equilibria in normal form games with compact action sets. This requires the use of the theory of correspondences. We summarise only the elements of this theory that are required for the proof of the existence of Nash equilibrium. For elaborate treatments of the theory of correspondences I refer to Hildenbrand (1974), Klein and Thompson (1984) and Aubin and Frankowska (1990).

Throughout we let X be a nonempty, convex and compact set in a finite dimensional Euclidean space. Furthermore, we assume that $\Phi: X \rightarrow 2^X$ is a set-valued correspondence on X such that $\Phi(x) \subset X$ is a compact and nonempty set, i.e., Φ is said to be nonempty- and compact-valued. We note that the next definitions are explicitly restricted for the type of correspondence considered here.

The correspondence Φ is called *upper hemi-continuous* (UHC) if for every sequence $x_n \rightarrow x \in X$ and every sequence y_n in X with $y_n \in \Phi(x_n)$ for every $n \in \mathbb{N}$, it holds that (y_n) has a convergent subsequence that converges to some $y \in \Phi(x)$.

Equivalently, the UHC property can be understood that the graph of the correspondence Φ is a closed set. Hence, $\{(x, y) \mid y \in \Phi(x)\} \subset X^2 = X \times X$ is a closed set in the corresponding Euclidean topology on X^2 .

Without proof we state the following well-known fixed-point theorem for correspondences in Euclidean spaces. This theorem was seminally stated and proven by Kakutani (1941). Kakutani's Theorem is used widely throughout game theory to show the existence of equilibria.

Theorem 1.6 (Kakutani's Fixed Point Theorem)

Let X be some compact, convex, and nonempty set of some finite dimensional Euclidean space. Let $\Phi: X \rightarrow 2^X$ be an upper hemi-continuous correspondence such that $\Phi(x) \subset X$ is compact, convex and nonempty. Then Φ admits a fixed point $x^ \in X$ such that $x^* \in \Phi(x^*)$.*



The aim for the proof of Theorem 1.3 is to show that the best-response correspondence $B: \mathbf{A} \rightarrow 2^{\mathbf{A}}$ satisfies Kakutani's Fixed Point Theorem 1.6. For that purpose we use the familiar notation \mathbf{A} instead of X and B instead of Φ . The proof of Theorem 1.3 proceeds through a number of intermediary steps that are formulated as lemmas.

Lemma 1.1

The set \mathbf{A} is a nonempty, compact and convex set in a finite dimensional Euclidean space.



Proof Indeed, since for every player $i \in N$ her action set A_i is a subset of some d_i -dimensional Euclidean space, it follows that the profile space $\mathbf{A} = \prod_{i \in N} A_i$ is a $d_1 \times \dots \times d_n$ -dimensional Euclidean space. Furthermore, since for every player $i \in N$ the action set A_i is nonempty, compact and convex, it follows immediately that $\mathbf{A} = \prod_{i \in N} A_i$ is nonempty, compact and convex as well. \square

Lemma 1.2

For every profile $a \in \mathbf{A}$ the corresponding best-response set $B(a) \subset \mathbf{A}$ is nonempty, convex and compact.



Proof Let $a \in \mathbf{A}$ be some profile.

First we note that $B(a) \neq \emptyset$. Indeed, for every $i \in N$ by continuity of π_i and A_i being nonempty and compact, it follows from Weierstrass Theorem—which states that a continuous function assumes maxima in a nonempty, compact set in a Euclidean space—that $B_i(a_{-i}) \neq \emptyset$. Hence, $B(a) = \prod_{i \in N} B_i(a_{-i}) \subset \mathbf{A}$ is nonempty as

well.

To show the convexity of $B(a)$, let $i \in N$ be some player. Take $a'_i, a''_i \in B(a_{-i})$ and denote

$$M_i = \max_{b_i \in A_i} \pi_i(b_i, a_{-i}) = \pi_i(a'_i, a_{-i}) = \pi_i(a''_i, a_{-i}).$$

Let $\lambda \in [0, 1]$. Then by convexity of A_i it follows that $\hat{a}_i = \lambda a'_i + (1 - \lambda)a''_i \in A_i$ and we have by the quasi-concavity of π_i in a_i that

$$\pi_i(\hat{a}_i, a_{-i}) \geq \lambda \pi_i(a'_i, a_{-i}) + (1 - \lambda) \pi_i(a''_i, a_{-i}) = M_i.$$

Hence, $\hat{a}_i = \lambda a'_i + (1 - \lambda)a''_i \in B_i(a_{-i})$. This shows that $B_i(a_{-i})$ is a convex set for every $i \in N$ and, therefore, that $B(a) = \prod_{i \in N} B_i(a_{-i})$ is a convex set as well.

To show that $B_i(a_{-i})$ is a closed set, let $(a_i^n)_{n \in \mathbb{N}}$ be some sequence in $B_i(a_{-i})$. Hence, $\pi_i(a_i^n, a_{-i}) = M_i$ for all $n \in \mathbb{N}$. Furthermore, since A_i is compact, without loss of generality we may assume that (a_i^n) is convergent with $a_i^n \rightarrow \tilde{a}_i \in A_i$. By continuity of π_i it is immediately clear that $\pi_i(\tilde{a}_i, a_{-i}) = M_i$. Hence, $\tilde{a}_i \in B_i(a_{-i})$, which implies that $B_i(a_{-i})$ is closed. Since $i \in N$ is arbitrary, it follows that $B(a)$ is a closed subset of \mathbf{A} . By compactness of \mathbf{A} it can now be concluded that $B(a)$ is indeed a compact set. \square

Lemma 1.3

The correspondence $B: \mathbf{A} \rightarrow 2^{\mathbf{A}}$ is UHC.



Proof It suffices to show that for every $i \in N$ the best response correspondence $B_i: \mathbf{A} \rightarrow 2^{A_i}$ is UHC. For that purpose let $(a_{-i}^n)_{n \in \mathbb{N}}$ be a sequence in $\prod_{j \neq i} A_j$ such that $a_{-i}^n \rightarrow \tilde{a}_{-i} \in \prod_{j \neq i} A_j$. Next, let $a_i^n \in B_i(a_{-i}^n)$ be a corresponding sequence of best responses to this sequence of action profiles. Hence, for any $b_i \in A_i$,

$$\pi_i(a_i^n, a_{-i}^n) \geq \pi_i(b_i, a_{-i}^n). \quad (1.13)$$

From the compactness of A_i , it can be assumed without loss of generality that $a_i^n \rightarrow \tilde{a}_i \in A_i$. By the continuity of π_i and by taking simultaneous limits in inequality (1.13), we immediately conclude that $\pi_i(\tilde{a}_i, \tilde{a}_{-i}) \geq \pi_i(b_i, \tilde{a}_{-i})$ for any $b_i \in A_i$. Hence, we conclude that $\tilde{a}_i \in B_i(a_{-i})$, showing that B_i is indeed UHC. \square

From the three lemmas above we conclude that Kakutani's Fixed Point Theorem 1.6 indeed applies to the best response correspondence B . Therefore, B has a fixed point, which is actually a Nash equilibrium of the game Γ . This completes the proof of Theorem 1.3.

1.5.3 Proof of Theorem 1.4

Let $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ be some normal form game. We now show the validity of the three assertions in Theorem 1.4 regarding the game Γ .

Proof of assertion 1.4 (a)

Assume that Γ admits at least one Nash equilibrium and that $\text{NE}(\Gamma) \neq \emptyset$. Define the function $F: \mathbf{A} \times \text{NE}(\Gamma) \rightarrow \{0, 1, \dots, n\}$ by

$$F(a, \hat{a}) = \#\{i \in N \mid a_i \neq \hat{a}_i\} \quad (1.14)$$

as the number of players that in $a \in \mathbf{A}$ do not select the Nash equilibrium $\hat{a} \in \text{NE}(\Gamma)$. We remark that F is a well-defined function, since $\text{NE}(\Gamma) \neq \emptyset$, that assumes a finite number of values and that attains a maximum and a minimum, even if \mathbf{A} and $\text{NE}(\Gamma)$ are infinite.

We introduce the potential function $\Psi: \mathbf{A} \rightarrow \{-n, \dots, -1, 0\}$ for every $a \in \mathbf{A}$ by

$$\Psi(a) = - \min_{\hat{a} \in \text{NE}(\Gamma)} F(a, \hat{a}) \quad (1.15)$$

Note first that the potential function only attains non-positive integer values only.

Denote by $\Gamma_\Psi = \langle N, \mathbf{A}, \Psi \rangle$ the corresponding potential form game for this potential function. Now we claim that it holds that $a \in \text{E}(\Gamma_\Psi)$ if and only if $\Psi(a) = 0$ if and only if $a \in \text{NE}(\Gamma)$.

First, we remark that if $\tilde{a} \in \text{NE}(\Gamma)$, then by definition $\Psi(\tilde{a}) = 0$. Second, if for $a \in \mathbf{A}$ we have that $\Psi(a) = 0$, then again by definition $a \in \text{NE}(\Gamma)$.

This implies that $\Psi(\hat{a}) = 0$ if and only if $\hat{a} \in \text{NE}(\Gamma)$.

Furthermore, if $\Psi(\hat{a}) = 0$, then obviously \hat{a} is a Ψ -maximiser, which implies that $\hat{a} \in \text{E}(\Gamma_\Psi)$.

Next, suppose that for $a \in \mathbf{A}$ it holds that $\Psi(a) < 0$. Then $\Psi(a) \leq -1$. Take $\hat{a} \in \text{NE}(\Gamma)$ such that $\Psi(a) = -F(a, \hat{a}) \leq -1$. Therefore, there is at least one player that selects in a an action different from her Nash equilibrium action in \hat{a} . Thus, we may select $i \in N$ such that $a_i \neq \hat{a}_i$. Therefore,

$$F(a, \hat{a}) = F((\hat{a}_i, a_{-i}), \hat{a}) + 1 > F((\hat{a}_i, a_{-i}), \hat{a}).$$

We now conclude that

$$\Psi(a) = -F(a, \hat{a}) < -F((\hat{a}_i, a_{-i}), \hat{a}) \leq \Psi(\hat{a}_i, a_{-i}).$$

This shows that $a \notin \text{E}(\Gamma_\Psi)$.

This implies that $a \in \text{E}(\Gamma_\Psi)$ if and only if $\Psi(a) = 0$.

This shows the assertion.

Proof of assertion 1.4 (b)

Suppose that $\text{NE}(\Gamma) = \emptyset$ and let $j \in N$ be such that A_j is an infinite action set. Take a countable selection in A_j denoted by $\{a_j^n \mid n \in \mathbb{N}\} \subseteq A_j$ consisting of different actions.

Next, we define the potential function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ by

$$\Psi(a) = \begin{cases} n & \text{if } a_j = a_j^n \\ 0 & \text{otherwise} \end{cases} \quad (1.16)$$

Now we claim that $\text{E}(\Gamma_\Psi) = \emptyset$.

Indeed, for $a \in \mathbf{A}$ it holds that either $a_j = a_j^n$ for some $n \in \mathbb{N}$, or $\Psi(a) = 0$. In either case, player j can deviate from a by selecting $a_j^{n+1} \in A_j$ to improve her payoff in Γ_Ψ to $\Psi(a_{-j}, a_j^{n+1}) = n + 1 > \Psi(a)$. Hence, a is

definitely not an equilibrium of Γ_Ψ .

Proof of assertion 1.4 (c)

Since Γ is a Nash potential game for a continuous potential function, we can take a continuous potential function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ such that $E(\Gamma_\Psi) = NE(\Gamma)$.

Since Ψ is continuous on the compact profile space \mathbf{A} , it follows from the Weierstrass Theorem that Ψ attains a maximum. Let $\hat{a} \in \mathbf{A}$ be a maximiser of Ψ .

Obviously, the maximiser \hat{a} is an equilibrium of Γ_Ψ . Therefore, $\hat{a} \in E(\Gamma_\Psi) = NE(\Gamma)$. Thus, \hat{a} is a Nash equilibrium of Γ .

1.5.4 Proof of Theorem 1.5

Let $\Gamma = \langle N, \mathbf{A}, \pi \rangle$ be a Euclidean normal form game with $NE(\Gamma) \neq \emptyset$. From Theorem 1.2 it then follows that $NE(\Gamma) \subset \mathbf{A}$ is a closed set in the corresponding Euclidean topology.

We define a function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ by

$$\Psi(a) = - \inf_{\hat{a} \in NE(\Gamma)} \|a - \hat{a}\| \quad (1.17)$$

We now prove the main assertion that Ψ is a continuous potential function such that $E(\Gamma_\Psi) = NE(\Gamma) \neq \emptyset$. We conduct this proof through three lemmas.

Lemma 1.4

The function $\Psi: \mathbf{A} \rightarrow \mathbb{R}$ is continuous on \mathbf{A} and it can be rewritten as

$$\Psi(a) = - \min_{\hat{a} \in NE(\Gamma)} \|a - \hat{a}\| \quad (1.18)$$

Proof First we show (1.18). Define for every $a \in \mathbf{A}$:

$$G(a) = \inf_{a^* \in NE(\Gamma)} \|a - a^*\|.$$

First, we remark that for every $a \in \mathbf{A}$ the value of $G(a)$ is finite due to $NE(\Gamma)$ being a closed, non-empty subset of \mathbf{A} .

If $a \in NE(\Gamma)$, it follows that $G(a) = 0 = \min_{a^* \in NE(\Gamma)} \|a - a^*\|$.

Next, let $a \notin NE(\Gamma)$. Then for every $n \in \mathbb{N}$ by definition of $G(a)$ there exists some $a_n^* \in NE(\Gamma)$ such that

$$\|a - a_n^*\| \leq G(a) + \frac{1}{n}.$$

We show that $(a_n^*)_{n \in \mathbb{N}}$ is a bounded sequence. Indeed,

$$\begin{aligned} \|a_n^*\| &\leq \|a\| + \|a_n^* - a\| \leq \|a\| + G(a) + \frac{1}{n} \\ &\leq \|a\| + G(a) + 1. \end{aligned}$$

Since \mathbf{A} is a closed set in a finite dimensional Euclidean vector space, $(a_n^*)_{n \in \mathbb{N}}$ has a convergent subsequence. Assume without loss of generality that $a_n^* \rightarrow \tilde{a} \in \mathbf{A}$. Since $(a_n^*)_{n \in \mathbb{N}} \subset NE(\Gamma)$ and $NE(\Gamma)$ is a closed set, it

follows that $\tilde{a} \in \text{NE}(\Gamma)$. Furthermore, $\|a - \tilde{a}\| = G(a)$. But this implies then that the infimum has been attained in \tilde{a} , implying that $G(a) = \min_{a^* \in \text{NE}(\Gamma)} \|a - a^*\|$.

Finally, continuity now easily follows from the property that G is a continuous function as defined. \square

Lemma 1.5

Every $\hat{a} \in \text{NE}(\Gamma)$ is a global maximiser of the function Ψ as defined and, therefore, $\hat{a} \in \text{E}(\Gamma_\Psi)$. 


Proof First, note that if $a \in \text{NE}(\Gamma)$, then it follows that $\Psi(a) = 0$.

Next, assume that $\Psi(a) = 0$. Then by Theorem 1.2 it holds that $a \in \overline{\text{NE}(\Gamma)} = \text{NE}(\Gamma)$.

Furthermore, if $a \notin \text{NE}(\Gamma)$, then obviously $\min_{\hat{a} \in \text{NE}(\Gamma)} \|a - \hat{a}\| > 0$. Thus, $\Psi(a) < 0$.

This implies the assertion that all Nash equilibria of Γ are Ψ -maximisers. \square

Lemma 1.6

Every $\tilde{a} \in \text{E}(\Gamma_\Psi)$ is a Nash equilibrium of Γ , i.e., $\tilde{a} \in \text{NE}(\Gamma)$. 

Proof The previous lemma showed that $\text{NE}(\Gamma) \subset \text{E}(\Gamma_\Psi)$. Hence, we need to show that every $a \in \text{E}(\Gamma_\Psi)$ is a Nash equilibrium of Γ .

Note that $a \in \text{E}(\Gamma_\Psi)$ if and only if for every $i \in N$: $b_i \in A_i$ implies that $\Psi(a) \geq \Psi(b_i, a_{-i})$. This is equivalent to

$$\min_{\hat{a} \in \text{NE}(\Gamma)} \|a - \hat{a}\| \leq \min_{\hat{a} \in \text{NE}(\Gamma)} \|(b_i, a_{-i}) - \hat{a}\|.$$

Let $\tilde{a} \in \arg \min_{\hat{a} \in \text{NE}(\Gamma)} \|a - \hat{a}\|$. Then

$$\|a - \tilde{a}\| \leq \min_{\hat{a} \in \text{NE}(\Gamma)} \|(b_i, a_{-i}) - \hat{a}\| \leq \|(b_i, a_{-i}) - \tilde{a}\|.$$

Therefore, $\|a - \tilde{a}\|^2 \leq \|(b_i, a_{-i}) - \tilde{a}\|^2$, implying that $(a_i - \tilde{a}_i)^2 \leq (b_i - \tilde{a}_i)^2$. We conclude therefore that $|a_i - \tilde{a}_i| \leq |b_i - \tilde{a}_i|$. Since, $b_i \in A_i$ is arbitrary, we conclude that $a_i = \tilde{a}_i$ for all $i \in N$. Hence, $a = \tilde{a}$.

This shows the assertion. \square

The three lemmas stated above show the assertion of Theorem 1.5.

Bibliography

- Anagnostides, Ioannis et al. (2025). “Convergence of Regret Matching in Potential Games and Constrained Optimization”. In: *arXiv* 2510.17067.
- Apostol, Tom M (1974). *Mathematical Analysis*. Reading, MA: Addison-Wesley.
- Aubin, Jean Pierre and H el ene Frankowska (1990). *Set-Valued Analysis*. Modern Birkh user Classics. Boston, MA: Birkh user.
- Axler, Sheldon (2020). *Measure, Integration & Real Analysis*. Graduate Texts in Mathematics 282. Berlin, Germany: Springer Verlag.
- (2024). *Linear Algebra Done Right*. Fourth edition. Undergraduate Text in Mathematics. Berlin, Germany: Springer Verlag.
- Beckmann, Martin, C. B. McGuire, and Christopher B. Winsten (1956). *Studies in the Economics of Transportation*. Vol. 12. Cowles Commission Monograph. First implicit use of a convex potential functional in traffic-assignment equilibrium; predates Rosenthal (1973). New Haven: Yale University Press.
- Bonanno, Giacomo (2018). *Game Theory: An open access textbook*. 2nd. Davis, CA: University of California Press.
- Candogan, Ozan, Ishai Menache, et al. (2011). “Flows and Decompositions of Games: Harmonic and Potential Games”. In: *Mathematics of Operations Research* 36.3, pp. 474–503.
- Candogan, Ozan, Asuman Ozdaglar, and Pablo A. Parrilo (May 2013). “Near-Potential Games: Geometry and Dynamics”. In: *ACM Transactions on Economics and Computation* 1.2, pp. 1–32. DOI: [10.1145/2465769.2465776](https://doi.org/10.1145/2465769.2465776).
- Cao, Zhigang, Zhibin Tan, and Jinchuan Zhou (2024). “Beyond Global Maximizers: Unveiling the Full Potential of Potential Games”. In: *SSRN Working Paper* 4704774.
- Cheng, Daizhan (2014). “On finite potential games”. In: *Automatica* 50, pp. 1793–1801.
- Cournot, Augustin (1838). *Recherches sur les principes math ematiques de la th orie des richesses*. Paris, France: Hachette.
- Debreu, G rard (1952). “A Social Equilibrium Existence Theorem”. In: *Proceedings of the National Academy of Sciences* 38, pp. 886–893.
- Fabrikant, Alex, Christos H. Papadimitriou, and Kunal Talwar (2004). “The Complexity of Pure Nash Equilibria”. In: *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC '04)*. ACM, pp. 604–612. DOI: [10.1145/1007352.1007445](https://doi.org/10.1145/1007352.1007445).
- Fan, Ky (1952). “Fixed Points and Minimax Theorems in Locally Convex Topological Linear Spaces”. In: *Proceedings of the National Academy of Sciences* 38, pp. 121–126.
- Folmer, H and P von Mouche (2004). “On a less known Nash Equilibrium Uniqueness Result”. In: *Journal of Mathematical Sociology* 28, pp. 67–80.
- Fudenberg, Drew and David K. Levine (1993). “Self-confirming Equilibrium”. In: *Econometrica* 61, pp. 523–545.
- Fudenberg, Drew and Jean Tirole (1991). *Game Theory*. Cambridge, MA: MIT Press.
- Glicksberg, Isidor L (1952). “A Further generalization of the Kakutani Fixed Point theorem with Applications to Nash Equilibrium Points”. In: *Proceedings of the National Academy of Sciences* 38, pp. 170–174.
- Green, George (1828). *An Essay on the Application of Mathematical Analysis to the Theories of Electricity and Magnetism*. Locus classicus for the term “potential function” in mathematical physics. Nottingham: T. Wheelhouse.
- Hargreaves-Heap, Shaun and Yanis Varoufakis (2004). *Game Theory: A Critical Introduction*. 2nd. London, UK: Routledge.
- Hart, Sergiu and Andreu Mas-Colell (1989). “Potential, Value, and Consistency”. In: *Econometrica* 57, pp. 589–614.
- Hildenbrand, Werner (1974). *Core and Equilibria of a Large Economy*. Princeton, New Jersey: Princeton University Press.
- Ismail, Mehmet S and Ronald Peeters (2024). “A connection between von Neumann-Morgenstern expected utility and symmetric potential games”. In: *Theory and Decision* 97, pp. 707–720.

- Jehle, Geoffrey A and Philip J Reny (2011). *Advanced Microeconomic Theory*. Third. Boston, Massachusetts: Addison Wesley Longman.
- Johnson, Samuel (1755). *A Dictionary of the English Language*. Two-volume dictionary; cited for the eighteenth-century senses of the word “potential”. London: W. Strahan.
- Kakutani, Shizuo (1941). “A Generalization of Brouwer’s Fixed Point Theorem”. In: *Duke Mathematical Journal* 8, pp. 457–459.
- Kellogg, Oliver Dimon (1929). *Foundations of Potential Theory*. Vol. 31. Grundlehren der mathematischen Wissenschaften. Canonical twentieth-century reference on classical potential theory. Berlin: Springer.
- Klein, Erwin and Anthony C. Thompson (1984). *Theory of Correspondences: Including Applications to Mathematical Economics*. Wiley & Sons.
- Lakatos, Imre (1978). *The Methodology of Scientific Research Programmes*. Cambridge, MA: Cambridge University Press.
- Marden, Jason R. (2012). “State Based Potential Games”. In: *Automatica* 48, pp. 3075–3088.
- Marden, Jason R., Gürdal Arslan, and Jeff S. Shamma (2009). “Cooperative Control and Potential Games”. In: *SIAM Journal on Control and Optimization* 47.3, pp. 1393–1407. DOI: [10.1137/060676386](https://doi.org/10.1137/060676386).
- Marden, Jason R. and Jeff S. Shamma (2015). “Game Theory and Distributed Control”. In: *Handbook of Game Theory with Economic Applications*. Ed. by H. Peyton Young and Shmuel Zamir. Vol. 4. Elsevier, pp. 861–899. DOI: [10.1016/B978-0-444-53766-9.00016-1](https://doi.org/10.1016/B978-0-444-53766-9.00016-1).
- Maschler, Michael, Eilon Solan, and Shmuel Zamir (2013). *Game Theory*. Cambridge, MA: Cambridge University Press.
- Monderer, Dov and Lloyd S Shapley (1996a). “Potential Games”. In: *Games and Economic Behavior* 14, pp. 124–143.
- (1996b). “Fictitious Play Property for Games with Identical Interests”. In: *Journal of Economic Theory* 68.1. Companion paper to Monderer–Shapley (1996); establishes convergence of fictitious play in games with identical interests, and by extension in exact potential games., pp. 258–265.
- Myerson, Roger B. (1991). *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press.
- Nash, John (1950). “Equilibrium Points in n -Person Games”. In: *Proceedings of the National Academy of Sciences* 36, pp. 48–49.
- (1951). “Non-cooperative Games”. In: *Annals of Mathematics* 54.2, pp. 286–295.
- Neumann, J. von and O. Morgenstern (1947). *Theory of Games and Economic Behavior*. 3rd. Princeton, NJ: Princeton University Press.
- Neyman, Abram (1997). “Correlated Equilibrium and Potential Games”. In: *International Journal of Game Theory* 26, pp. 223–227.
- Nocke, Volker and Nicolas Schutz (2025). “Oligopoly, Complementarities, and Transformed Potentials”. In: *CRC TR 224 Discussion Paper* 644.
- Osborne, Martin J (2004). *An Introduction to Game Theory*. New York, NY: Oxford University Press.
- Osborne, Martin J and Ariel Rubinstein (1994). *A Course in Game Theory*. Cambridge, Massachusetts: MIT Press.
- Owen, Guillermo (2013). *Game Theory*. 4th ed. Bingley, UK: Emerald Publishing Group.
- Rockafellar, R Tyrrell (1970). *Convex Analysis*. Princeton, NJ: Princeton University Press.
- Rosenthal, Robert W. (1973). “A Class of Games Possessing Pure-Strategy Nash Equilibria”. In: *International Journal of Game Theory* 2.1, pp. 65–67. DOI: [10.1007/BF01737559](https://doi.org/10.1007/BF01737559).
- Roughgarden, Tim (2006). “Potential Functions and the Inefficiency of Equilibria”. In: *Proceedings of the International Congress of Mathematicians*. Madrid, Spain: European Mathematical Society, pp. 1071–1094.
- (2016). *Twenty Lectures on Algorithmic Game Theory*. Textbook exposition of algorithmic and computational game theory; includes a chapter-length tutorial on potential games, price of anarchy, and convergence of no-regret dynamics. Cambridge: Cambridge University Press.
- Shapley, Lloyd S (1953). “A Value for n -Person Games”. In: *Contributions to the Theory of Games*. Ed. by R.D. Luce and A.W. Tucker. Vol. II. Princeton, NJ: Princeton University Press.
- Slade, Margaret E. (1994). “What Does an Oligopoly Maximize?” In: *The Journal of Industrial Economics* 42.1, pp. 45–61.

- Strang, Gilbert (2021). *Introduction to Linear Algebra*. 5th. Cambridge, MA: Wellesley-Cambridge Press.
- Swenson, Brian, Ryan Murray, and Soumya Kar (2020). “Regular Potential Games”. In: *Operations Research*.
- Ui, Takashi (2000). “A Shapley Value Representation of Potential Games”. In: *Games and Economic Behavior* 31, pp. 121–135.
- (2001). “Robust Equilibria of Potential Games”. In: *Econometrica* 69.5, pp. 1373–1380.
- Vega-Redondo, Fernando (2003). *Economics and the Theory of Games*. New York, NY: Cambridge University Press.
- Voorneveld, Mark (1999). “Potential Games and Interactive Decisions with Multi-Criteria”. PhD thesis. Tilburg, The Netherlands: Tilburg University.
- Weirich, Paul (1998). *Equilibrium and Rationality: Game Theory Revised by Decision Rules*. Cambridge, United Kingdom: Cambridge University Press.