Tilburg University

Game theory

van Damme, E.E.C.

Published in:
International Encyclopedia of the Social & Behavioral Sciences

DOI:
[10.1016/B978-0-08-097086-8.71048-8](https://doi.org/10.1016/B978-0-08-097086-8.71048-8)

Publication date:
2015

Document Version
Publisher's PDF, also known as Version of record

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
van Damme, E. E. C. (2015). Game theory: Noncooperative games. In J. Wright (Ed.), *International Encyclopedia of the Social & Behavioral Sciences* (2nd ed., Vol. 9, pp. 582-591). Elsevier.
<https://doi.org/10.1016/B978-0-08-097086-8.71048-8>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Copyright © 2015 Elsevier Ltd. unless otherwise stated. All rights reserved.
Elsevier

Game Theory: Noncooperative Games

Eric van Damme, Tilburg University, Tilburg, The Netherlands

© 2015 Elsevier Ltd. All rights reserved.

Abstract

We describe noncooperative game models and discuss game theoretic solution concepts. Some applications are also noted. Conventional theory focuses on the question 'how will rational players play?', and has the Nash equilibrium at its core. We discuss this concept and its interpretations, as well as refinements (perfect and stable equilibria) and relaxations (rationalizability and correlated equilibria). Motivated by experiments that show systematic theory violations, behavioral game theory aims to integrate insights from psychology to get better answers to the question 'how do humans play?'. We provide an overview of the observed regularities and briefly sketch (beginnings of) theories of boundedly rational play.

Introduction

Games are mathematical models of interactive decision situations, i.e., situations in which multiple decision makers, each one with its own objectives, jointly determine the outcome. Game theory aims to predict what players will do in such situations and what outcomes will result. The theory has been applied in economics, other social sciences, biology and computer science, among others. Aumann (1987) presents an overview of how the field developed in the twentieth century. The 3-volume *Handbook of Game Theory with Economic Applications* (Aumann and Hart, 1992/1994/2002) provides a fairly complete overview of rationalistic game theory in almost 2400 pages. Excellent textbooks at the graduate level are Myerson (1991) and Osborne and Rubinstein (1994). For more information and detailed references, the reader is advised to consult any of these sources. Behavioral game theory is more recent, less established and developing more quickly; Camerer (2003) provides a good starting point for this branch.

The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel has been awarded to game theorists four times. In 1994, John Harsanyi, John Nash and Reinhard Selten shared the Prize for developing equilibrium theory in noncooperative games; in 2005 Robert Aumann and Thomas Schelling shared it for enhancing our understanding of conflict and cooperation by means of game theoretic analysis; in 2007 Leonid Hurwicz, Eric Maskin, and Roger Myerson were praised for developing the theory of mechanism design, while in 2012 Alvin Roth and Lloyd Shapley received the Prize for the theory of stable allocations and the practice of market design. Researchers from closely related fields were honored in 1996 (James Mirrlees and William Vickrey; incentives under asymmetric information), in 2002 (Daniel Kahneman, Vernon L. Smith; behavioral economics and experimental economics), and in 2009 (Elinor Ostrom (who shared the Prize with Oliver Williamson); institutional economics). Excellent information on the contributions of the Prize winners is available on the official web site of the Nobel Prize, http://www.nobelprize.org/nobel_prizes/economics/laureates/index.html. In 2014, the Prize was awarded to Jean Tirole, who also made important contributions to noncooperative game theory, among others together with Drew Fudenberg and Eric Maskin.

The game theory literature distinguishes two main classes of models: cooperative games and noncooperative games. The terminology, which suggests that in one case the players

cooperate and in the other they do not, is misleading. The difference is not in what players want, but rather in what they are allowed to do. Traditional game theory assumes that players are rational and strive to maximize their utility; in this respect, there is no difference between the two models. In noncooperative theory, however, it is assumed that the model is complete and that the players are bound by its rules. In particular, contracts or commitments are binding only if the formal rules explicitly allow this. By contrast, in cooperative theory, players are free to negotiate, form coalitions and possibly make side payments, and are assumed to have access to a costless external mechanism that enforces agreements. Noncooperative theory assumes such an external mechanism is absent, hence, focuses on self-enforcing agreements.

The distinction was coined by John Nash in his Ph.D. thesis (Nash, 1950), which also introduced the fundamental solution concept for noncooperative games (the Nash equilibrium concept). In the path-breaking Von Neumann and Morgenstern (1944), the founders of game theory had developed two distinct theories: one for two-person games in which the players have strictly opposite interests (two-person zero-sum games) and another for n -person games in which the players can form coalitions and make side payments. They had argued that, as soon as there are more than two players, choosing an ally and forming a coalition becomes crucial, with side payments being key in stabilizing cooperation; hence, that two-person zero-sum games, in which such possibilities are irrelevant, are the exception. Consequently, they assumed that a mechanism enforcing coalitions and contracts was available and focused on cooperative theory. Nash extended Von Neumann and Morgenstern's two-person zero-sum theory and developed the general noncooperative theory. In this theory, each player acts independently, without collaborating with any of the others, however, by making full use of all the possibilities for cooperation that the game allows.

A model should be rich enough to allow for the relevant possibilities, but also simple enough to allow in-depth analysis and yield insight. Nash argued that noncooperative models are more fundamental, as it should always be possible to model coalition negotiations as formal moves in a noncooperative game. Although correct, the resulting model may be too complicated and the attention for details may blur the general picture, hence, each type of model has its advantages. Possibilities for cooperation can be modeled either as part of the game, or as part of the solution concept. In this contribution,

we limit ourselves to the former approach. As we will see, there exist various deep links between the two approaches. Cooperative game theory is surveyed in William Thomson's contribution to this Encyclopedia.

The remainder of the material is structured as follows. We first describe the two main classes of noncooperative models (the extensive form and the strategic (or normal) form), and the concept of strategy that allows the reduction of one model to the other. Next, we turn to solution concepts that are based on the assumption that players are perfectly rational and have full understanding of the game. We discuss Nash equilibrium, some of its drawbacks, as well as extensions (correlated equilibria) and refinements of it, such as perfect and stable equilibria. The experimental literature has shown that human players may deviate from perfect rationality in systematic ways; hence, we next discuss recent results from the behavioral game theory literature, which aims to construct models of thinking and learning that are descriptively more accurate. We close by briefly discussing some applications, including the link between cooperative and noncooperative theory.

Noncooperative Game Models

An 'extensive game' is a very detailed model of a conflict situation; it specifies which players are involved and how the game evolves over time: which player moves when, what information does the player then have, what can he do, what are the possible consequences of his actions, and how do the players evaluate the outcomes? Von Neumann and Morgenstern (1944) already provided a set-theoretic description of this model, but Kuhn (1953) provided a graph theoretic formulation that is easier to work with and that has become the standard.

A special case is a game with 'perfect information,' in which the moves are sequential and each player, whenever he has to move, is fully informed about everything that has happened before. Chess is a game with perfect information. Figure 1 gives a very simple example: player 1 (P_1) moves first and chooses between terminating the game, O , with payoff 1 to P_1 and X to P_2 or giving the move to P_2 (action I), who then determines whether each player gets 0, or whether P_1 gets 2, with P_2 getting Y . We will return to this game below. Most of the literature has restricted attention to games with 'perfect recall' in which a player never forgets what he knew or what he has done

Figure 1 $X > Y > 0$

before. Bridge, when modeled as a two-player game, has imperfect recall: if EW is the defending team, then W does not know (recall) the cards of E.

By means of the concept of strategy, already introduced in early work of John von Neumann, an extensive form game can be reduced to its strategic form. A 'strategy' for a player is a full plan of action for how to play the game, i.e., it specifies a unique action for each decision point of this player and each piece of information that this player might then have. Denote by N the set of players in the game and by S_i the set of all strategies for player $i \in N$. Assume that players evaluate outcomes by Von Neumann Morgenstern utility functions and let u_i be the utility function of player i . An n -tuple of strategies $s = (s_1, \dots, s_n)$, one for each player, determines a probability distribution over the outcomes and, hence, implies a unique expected utility $u_i(s)$ for each player. Von Neumann argued that the 'strategic form' of the game, the tuple $\langle N, S_1, \dots, S_n, u_1, \dots, u_n \rangle$ specifying players, strategies and payoffs, contains all the information that is needed for rational players to determine what to do. Consequently, it would suffice to develop theory for 'strategic games.' The literature has debated whether the details of the extensive form are indeed irrelevant (see Section Rationality and Equilibrium). Table 1 gives the strategic form of the game from Figure 1: the rows are the strategies of P_1 , the columns the strategies of P_2 , and in each cell, the first number is the payoff (i.e., the utility) to P_1 and the second the payoff to P_2 ; a convention that will be followed throughout. Note that in a strategic game each player only moves once, with players choosing strategies simultaneously.

The interpretation of a 'strategic game' $\langle N, S, u \rangle$ is that the data of the game are common knowledge: all players are fully informed who the players are, what strategies each of them has available and how all players evaluate the possible outcomes. Harsanyi (1968) showed how incomplete information can be incorporated into the model. In Harsanyi's model, the game starts with a chance move that distributes private information to each player. It is then common knowledge what pieces of information each player might have and with what probabilities, but the exact piece of information (also called the player's type) is only known to the player himself. Formally, a 'Bayesian game' is a tuple $\langle N, T, S, p, u \rangle$, where N is the set of players, $T = T_1 \times \dots \times T_n$ is the type space, p is a probability distribution on T , and u is the vector of utility functions, where each player i 's utility may generally depend on all types and all actions taken, $u_i = u_i(t, s)$.

In a Bayesian game, information is distributed asymmetrically: P_i is informed about his type t_i and can choose his action $s_i \in S_i$ on the basis of it; P_j only knows his own type t_j and on this basis forms beliefs $p(t_i, t_j)$ about what P_i might know and might do, $s_i(t_i)$. As an example of a Bayesian game, think of a sealed bid auction: each bidder knows what the auctioned object is worth to him, but has only imprecise information

Table 1 A strategic game

	L	R
O	1, X	1, X
I	0, 0	2, Y

about the value of others. Also *see* the article on Auctions. A special class of Bayesian games, ‘signaling games,’ has proved fertile for theory development. In such a game, there is only one player with private information, but this player moves first; he has to decide how much information to reveal, while his opponents have to figure out what information his action might be signaling (*see* the article Information, Economics of).

It should be noted that above we referred to common knowledge in a loose sense: all players knowing something is different from that being common knowledge. The latter also implies that all players know that all players know it, and that all players know that all players know that all players know it, etc; *see* Aumann (1976) for a formal definition. Rubinstein (1989) shows that games with almost common knowledge are very different than games with common knowledge.

In the above discussion, we restricted attention to pure strategies; however, there is also the possibility of randomization. In a strategic form game, a ‘mixed strategy’ of P_i is a probability distribution σ_i over this player’s set of pure strategies S_i . (We will write $\Delta(S_i)$ for the set of all such probability distributions.) It can be interpreted either as an act of deliberate randomization of this player, or as an expression of the uncertainty that the other players face about what P_i is going to do. In the latter case, σ_i represents the common beliefs held by the opponents of P_i . (The assumption that they have the same beliefs makes sense if they have the same information; in a Bayesian game, σ_i are the *ex-ante* beliefs, before the own type is known.)

In an extensive game, one can distinguish two types of randomization. When using a mixed strategy, the player randomizes over his pure strategies before the game starts. If the player randomizes locally over his actions at each of his information sets, this is called a ‘behavior strategy.’ A mixed strategy always induces a behavior strategy, but the converse only holds in games with perfect recall (Kuhn, 1953); hence, in these games, the restriction to behavior strategies is without loss of generality. Bridge provides an illustration that, without perfect recall, one may be able to play better with mixed strategies: to play optimally, the defending team needs to perfectly coordinate the actions of its members without revealing too much information; this is possible in a mixed strategy, but not when the team members randomize independently.

Rationality and Equilibrium

Nash Equilibrium

The fundamental solution concept for noncooperative games was introduced by John Nash in his 1950 PhD thesis. The mathematical core of the thesis was published as Nash (1951), but the chapter ‘Motivation and Interpretation,’ was not published, which may have led to misunderstandings and may have delayed development of the field.

Let $G = \langle N, S, u \rangle$ be an n -person strategic game. A ‘Nash equilibrium’ of G is a strategy combination of σ (either mixed or pure) with the property that each player i is playing a best response against the strategies played by the others, hence $u_i(\sigma) = \max_{\tau_i} u_i(\tau_i, \sigma_{-i})$ for all $i \in N$. In other words, as long as the others do not deviate from their equilibrium strategies,

player i cannot improve his payoff by deviating from his equilibrium strategy. An equilibrium of an extensive game is defined similarly; it simply is an equilibrium of the associated strategic game. An equilibrium of a Bayesian game is frequently called a Bayesian Nash equilibrium: each player plays a best response against the strategies of the others, whatever his type might be.

Nash provided two justifications for his concept. The first interpretation is rationalistic; Nash equilibrium is an answer to the question: what would perfectly rational players do? If we assume that a theory of rational play produces a unique solution and if the players know the solution, then rational (payoff maximizing) players will conform to this solution only if it is a Nash equilibrium. Any other (single-valued) theory is self-defeating. While assuming rational players to know the solution and to make use of their knowledge seems fine, the assumptions of existence and uniqueness are crucial. In fact, as a game may have multiple Nash equilibria (*see* below), this rationalistic justification seems incomplete at best. How can a player predict what another player will do if there are multiple Nash equilibria? This question motivated a fruitful, long-term research project of two Nobel Prize winners, who ultimately showed that additional, but not undisputable, assumptions allow one to come up with an answer; Harsanyi and Selten (1988).

The second ‘mass action interpretation’ assumes the game to be repeated, with each time players being newly drawn from certain populations and with each player accumulating empirical information on the relative advantages of their own strategies as well as on how often the opponents play their strategies. If the frequencies with which the various pure strategies are used converge, then “the mixed strategies representing the average behavior in each of the populations form an equilibrium point” (Nash, 1950: p. 22). In other words, under certain assumptions, learning leads to Nash equilibrium. Note that for this second interpretation, uniqueness is irrelevant; the initial conditions may determine at which equilibrium the process ends up. A large literature has investigated various types of learning processes and under which conditions convergence to equilibrium is indeed obtained; *see* Fudenberg and Levine (1998, 2009) and Young (2004) for overviews.

A third interpretation originates in the biological literature; *see* Smith (1982). Strategies are assumed to be randomly matched against each other with the payoff $u_i(\sigma)$ representing the fitness (expected number of offspring) of strategy σ_i when the state of the system is σ . The fittest strategies grow fastest, hence, if the system converges, it must be to a Nash equilibrium. The literature has studied various evolutionary processes, among which the replicator equation; *see* Weibull (1995).

By using a fixed point theorem (such as Brouwer’s or Kakutani’s) one can show that any strategic game has a Nash equilibrium, provided one allows equilibria in mixed strategies. ‘Matching pennies’ (two players simultaneously choose H or T with P_1 winning both pennies if the choices match and P_2 winning otherwise), shows that not all games have equilibria in pure strategies: in this game each player has to randomize.

Mixed strategies and mixed strategy equilibria are frequently viewed as being problematic, as players do not have

strict incentives to play them: any pure strategy in the support of a mixed equilibrium strategy is a best response as well. Mixed strategies, however, can also be interpreted as beliefs. Harsanyi (1973) showed that mixed equilibria arise naturally as beliefs associated with pure equilibria of a Bayesian game in which the uncertainty that each player faces about the other players' payoffs is explicitly taken into account. If one allows for the payoffs to be slightly uncertain, but each player being perfectly informed about his own payoff, then we have a larger Bayesian game $G(\varepsilon)$ in which each player i can play a pure strategy as for each payoff realization a specific pure action is optimal. However, as i 's opponents do not know i 's payoff realization, they will be uncertain about what action P_i will actually play. Given an equilibrium σ of a generic strategic game G one can find an equilibrium $\sigma(\varepsilon)$ of the Bayesian game $G(\varepsilon)$ such that, in the limit, as the uncertainty vanishes, for each $i \in N$ the beliefs of i 's opponents associated with $\sigma(\varepsilon)$ converge to σ_i .

As noted, Nash's noncooperative model is a generalization of Von Neumann and Morgenstern's two-person zero-sum game. Indeed, Nash's equilibrium concept is a generalization of their minimax solution. The founders of game theory asked the question what is the highest payoff that a player can guarantee himself and they defined a minimax strategy as one that guarantees this value. For two-person zero-sum games, σ is a Nash equilibrium if and only if, for each player i , σ_i is a minimax strategy. If games are not strictly competitive, however, the two concepts differ; best responding against a player that pursues his own interests is different from optimally defending yourself against somebody that plays against you.

Table 2 lists three well-known games. In the Prisoners' Dilemma (Table 2(a)), (D, D) is the unique Nash equilibrium. This shows that an equilibrium may be (Pareto) inefficient; another outcome is preferred by both players. That constraints on cooperation can hurt players is unsurprising; they no doubt would agree on (C, C) , if they could sign binding contracts. 'Battle of the Sexes' (Table 2(b), with $X = 0$) is a game with two pure equilibria, and a mixed strategy one yielding each player $5/6$, hence, a game may have multiple equilibria. In stag hunt (Table 2(c), with $0 < X < 4$), (S, S) and (R, R) are Pareto ranked equilibria. Although (R, R) yields higher payoffs for both players, S is a safer strategy: it guarantees the payoff X , while, if only one player chooses R , he ends up with 0. In this game, there is a conflict between risk dominance and payoff dominance: if $2 < X < 4$, then (S, S) risk dominates (R, R) (Harsanyi and Selten, 1988). The strategic game from Table 1 (with $Y > 0$) shows that some Nash equilibria may be unstable. (O, L) is a Nash equilibrium, but, for P_2 , R is always at least as good

as L and sometimes it is strictly better; hence, R weakly dominates L . We conclude that Nash equilibria always exist, that there may be multiple (nonequivalent) equilibria, and that some equilibria may be unstable.

By means of a model similar to Harsanyi's, Carlsson and Van Damme (1993) have shown that common payoff uncertainty can eliminate equilibria, hence, can serve as an equilibrium selection device. Consider a game as in Table 2(c). If $X < 0$, then S is a strictly dominated strategy; R is strictly dominated if $X > 4$, while for intermediate values, there are multiple equilibria. Carlsson and Van Damme consider the situation where X can take any real value, with players facing uncertainty, but each player receiving a reasonably accurate independent signal X_i about the true value of X before making his decision. They coined the term 'global game' for the resulting Bayesian game. If P_i receives a very large signal ($X_i \gg 4$), he can be reasonably sure that R is dominated, hence he will play S . Similarly, each player will choose R if his signal is very negative. It is natural to focus on simple (switching point) equilibria of the global game in which each P_i plays R if and only if $X_i < x^*$, for some x^* ; indeed the authors show that under certain assumptions only such equilibria exist. In such a Bayesian equilibrium, a player must be indifferent when receiving the signal x^* . However, as each player believes that the events $X_1 > X_2$ and $X_1 < X_2$ are approximately equally likely, indifference can hold only if, $x^* = 2$. In the limit, as uncertainty vanishes, players, hence, coordinate on the risk dominant equilibrium of the stag hunt. For further discussion on global games, on the conditions under which these have unique limit equilibria, the relation with common knowledge, and their applications, the reader is referred to Morris and Shin (2003).

Rationalizability, Iterated Dominance and Correlated Equilibria

Nash equilibrium assumes that players optimize and have correct beliefs about their opponents. The concept of rationalizability (Bernheim, 1984; Pearce, 1984) keeps the first assumption, but relaxes the second. It can be obtained as outcome of an interactive process. In the first step, all beliefs are allowed and any player i can choose any best response to these. In the second step, for any $i \in N$, the beliefs of i 's opponents are only allowed to put positive weight on strategies that are best responses for P_i and any player j is only allowed to play strategies that are best responses to profiles of the resulting beliefs. The 'rationalizable strategies' are those that survive iterative application of this procedure. In games with a unique rationalizable outcome, weaker rationality assumptions suffice to obtain the outcome.

The above procedure is related to (but not fully equivalent with) the iterative elimination of strictly dominated strategies, where a pure strategy r_i of player i is 'strictly dominated' if there exists a mixed strategy σ_i such that $u_i(r_i, s_{-i}) < u_i(\sigma_i, s_{-i})$ for all s_{-i} . The difference is related to the question of whether beliefs about different players can be correlated or not. This issue does not arise in two-person games, for which the two procedures are equivalent. Note that any Nash equilibrium is rationalizable, however, Nash equilibria may vanish if weakly dominated strategies are eliminated (a pure strategy r_i is 'weakly

Table 2 Three strategic form games

<i>(a) Prisoners' dilemma</i>			<i>(b) Battle of the sexes</i>			<i>(c) Stag hunt</i>		
	<i>C</i>	<i>D</i>		<i>L</i>	<i>R</i>		<i>S</i>	<i>R</i>
C	3, 3	0, 4	U	5, 1	0, 0	S	X, X	X, 0
D	4, 0	1, 1	D	X, X	1, 5	R	0, X	4, 4

dominated' if there exists s_{-i} satisfying the above inequalities for some s_{-i} and corresponding weak inequalities for all s_{-i}).

Aumann (1974) introduced the concept of correlated equilibrium that generalizes Nash's concept to games in which communication between the players is possible. It assumes that players can conduct joint lotteries (play correlated strategies), but cannot make binding agreements or side payments. As an introduction, consider the game from Table 2(b) with $X = 0$ ('Battle of the Sexes'). If players can communicate, they can decide to throw a coin together and to play (U, L) when the outcome is H and (D, R) , when the outcome is T . This correlated strategy yields each player the payoff 3, which is a good compromise. Furthermore, the agreement is self-enforcing: whatever the outcome of the coin toss, one player has a strong incentive to abide by the agreement and to follow up on it, so that it is in the best interest of the other to do so as well. We have a correlated equilibrium: a correlated strategy in which no player has an incentive to deviate from.

It will be clear that any convex combination of Nash equilibria is a correlated equilibrium. However, we can do more. Consider again the game from Table 2(b), but now with $4 < X < 5$ and consider the scenario in which the players instruct a trustworthy mediator to randomize equally among the three cells of the matrix that have positive payoffs. Furthermore, they instruct him that, for each outcome of the lottery, P_1 shall only be informed about which row resulted while P_2 shall only be informed about the column. Viewing this information as a recommendation of what to play, one notices that, if one player always follows the recommendation, it is in the interest of the other player to do so as well. For example, if U is recommended, P_1 knows that P_2 has been recommended (and will play) R , hence, U yields the highest payoff. Similarly, if D is recommended, P_1 knows that P_2 will play L or R , each with probability $1/2$, hence, again following the recommendation is best. The entire scheme is self-enforcing, hence, a correlated equilibrium, but it is not a convex combination of Nash equilibria. Formally, a 'correlated equilibrium' is a correlated strategy $\sigma \in \Delta(S)$ such that, for each player i , if $\sigma_i(s_i) > 0$, then s_i is a best response against $\sigma_{-i}(s_{-i} | s_i)$; in words: any recommendation s_i that any P_i might receive is a best response given i 's beliefs after hearing s_i .

Equilibrium Refinements and Stable Equilibria

We now focus attention on games in extensive form. Zermelo (1913) already showed that, in theory, chess can be solved by a backward induction procedure: starting at the end of the game, one works backwards replacing each decision point with the outcome that is obtained if a local best reply is taken there. It is straightforward to extend this procedure to any game with perfect information. Note that this procedure assumes persistent rationality: whatever happened before, each player assumes that all players will act rationally from then on. This assumption seems appropriate, but it is noteworthy that Von Neumann and Morgenstern (1944) already criticized it.

In the game of Figure 1, backward induction produces the Nash equilibrium (I, R) . There is, however, a second equilibrium, (O, L) , that is not consistent with this procedure. In effect,

in this second equilibrium, P_2 threatens to choose L and, if believed, P_1 chooses O , so that the threat does not have to be executed. The question is whether the threat to play L really is credible. What is at issue here is that the strategic form of the game seems to assume that a player can commit himself to a strategy. However, in a noncooperative game such commitments are impossible, hence, when faced with the *fait accompli* that P_1 has chosen I , the best P_2 can do is to choose R . Strategies should not just be optimal at the beginning of the game, but also from each decision point onwards; in a noncooperative setting, only backward induction equilibria make sense. When a player is rational, the possibility to reoptimize should not lead him to deviate from his original strategy.

Following Selten (1965), a large literature on 'equilibrium refinements' has studied the question of how to generally eliminate Nash equilibria that rely on 'incredible threats.' Three related but conceptually different strands of literature can be distinguished.

The first line groups concepts that aim to extend the backward induction procedure beyond games with perfect information. Selten (1965) proposed 'subgame perfect equilibria': equilibria that induce Nash equilibria in all subgames; a subgame being a part of the game tree that constitutes a game of itself. Kreps and Wilson (1982) strengthened this idea and introduced 'sequential equilibria.' Such an equilibrium consists of a strategy profile σ , together with a system of beliefs μ that for each player specifies a probability distribution over the nodes in each information set. Two conditions are required:

1. Sequential rationality: at each information set, the player's strategy is optimal against the strategies of the others given the beliefs.
2. Consistency: the system of beliefs should be compatible with the strategy profile.

Various formalizations of consistency have been proposed. The main advantage of this framework is that it provides a natural language to discuss 'reasonableness' of beliefs and of the associated equilibria, hence, sequential equilibria can be further refined by imposing additional conditions on the beliefs. For example, in signaling games, one can insist that, upon observing action a_i , uninformed players assign beliefs $\mu(t_i, a_i) = 0$ to those types t_i of the informed player for which a_i is dominated or equilibrium dominated (see the article on Alternative Schools of Economic Thought).

In the second strand of literature, it is assumed that players will, with a small probability, make mistakes and it is required that equilibria be robust against this possibility. Hence, perfect rationality is viewed as a limiting case of slightly imperfect rationality. The seminal paper in this strand is Selten (1975), which introduced the concept of 'perfect equilibrium,' based on the idea that all players make independent mistakes at each of their information sets. When players tremble, all decision nodes are reached with positive probability, so that an optimizing player is forced to choose a best response everywhere. Consequently, in the perturbed game there cannot be incredible threats. An equilibrium is said to be perfect if it can be obtained as limit of Nash equilibria of a sequence of perturbed games when the mistakes vanish. As Kreps and Wilson (1982) define their main consistency requirement in terms of mistakes, there is a clear

link between the two strands: any perfect equilibrium is sequential. The perfectness concept can be refined by imposing further conditions on the mistakes. For example, Roger Myerson's proper equilibrium insists that more costly mistakes occur with much lower frequency.

Sequential equilibria and perfect equilibria rely essentially on the extensive game structure. As a result, two extensive games with the same strategic form may have different sequential or perfect equilibria. Kohlberg and Mertens (1986) argued that such dependency is undesirable: fully rational players are not misled by presentation details that are strategically irrelevant. Hence, they argued in favor of a solution that satisfies invariance, i.e., which only depends on the strategic form. The game from Figure 1 illustrates that the backward induction property may be uncovered in the strategic form: (I, R) , the only strategy pair to survive backward induction in the extensive game, is also the only one surviving interactive elimination of weakly dominated strategies in the strategic form. Hence, we have two different rationality principles that produce the same outcome. More generally, a proper equilibrium of a strategic game induces a sequential equilibrium in any extensive form game with that strategic form. Hence, it seems possible that 'robust' equilibrium outcomes can be identified in the strategic form.

Kohlberg and Mertens (1986) initiated the axiomatic approach to equilibrium refinement: they postulate several properties that a rational solution should satisfy and investigate whether a solution satisfying these properties exists. Examples of such properties are: (1) invariance (already discussed above), (2) consistency with one-person decision theory (admissibility), (3) independence of strategies that are dominated or that are suboptimal responses against the solution, and (4) a solution should remain whenever a game is embedded in a larger one (the small worlds property). Kohlberg and Mertens (1986) proposed to strengthen perfectness by insisting not just on stability against one particular sequence of trembles, but against all small trembles. As typically a single equilibrium will not have this property, they suggested looking at minimal closed and connected sets of equilibria that are stable in this sense. This initial attempt did not satisfy all properties that they considered desirable, but Mertens (1989) next proposed a concept that indeed satisfies all of them. The definition of Mertens stability is highly technical, insisting on certain homology properties of the best reply correspondence. Govindan and Wilson (2008) defined the related concept of metastable equilibria that is somewhat weaker than Mertens stability, but satisfies the same decision-theoretic properties.

We thus conclude that the question 'how to exclude Nash equilibria that rely on incredible threats?' has led to highly technical questions about the best reply correspondence of the game. Exactly why such sophisticated techniques appear necessary to solve such an intuitive question is still imperfectly understood. Nevertheless, that stability suffices for that purpose can be shown with a simple example, which also illustrates the concept of forward induction. Consider the game in which P_1 first chooses whether to take up an outside option yielding him payoff 2 or to play Battle of the Sexes (the game from Table 2(b) with $X = 0$). Taking up the outside option is part of the perfect equilibrium (OD, R) : if P_1 thinks that (D, R) will be

played in the subgame, he is better off taking his option. However, this outcome does not seem reasonable: P_1 not choosing the outside option and then playing D is strictly dominated. Being requested to play, it seems that P_2 should, therefore, conclude that P_1 will play U and should respond with L . Hence, only the outcome $(5, 1)$ seems reasonable. Indeed, this is the only stable equilibrium outcome.

This same outside option game may, however, illustrate that in case communication is possible (hence, when the basic solution concept is correlated equilibrium), we cannot insist on the solution to just depend on the strategic form. Exactly when the communication takes place may matter. If players can only communicate before the start of the game, P_1 can never be induced to play ID , hence, communication is immaterial and the outcome is $(5, 1)$. On the other hand, if players can communicate after P_1 has thrown away his option, then players can randomize between $(5, 1)$ and $(1, 5)$, hence, P_1 can be induced to give up his option. Extensive form correlated equilibria are different from strategic form correlated equilibria, and for good reasons; see Myerson (1991; Chapter 6), where one can also find some remarks on how to refine correlated equilibria.

Behavioral Game Theory

Conventional game theory, with its focus on the question 'how will rational players play?' frequently makes sharp predictions about the outcome or about how this outcome changes with a change in the data (comparative statics), hence, studying the theory's empirical relevance appears quite natural. Testing with field data, however, has its limitations, and although there are exceptions, serious experimental investigation of the descriptive relevance of rationality-based theory only started in the 1980s (Kagel and Roth, 1995). The first wave of experimental studies established that standard theory sometimes (or frequently, depending on one's viewpoint) provides poor predictions of how humans play, and that there are systematic patterns in the deviations, which has then led to revised theories (or at least models) of play incorporating these regularities. In the last 25 years, emphasis has thus shifted to the question 'how do humans play non-cooperative games?' leading to a strong interaction between theory and empirical work.

This section provides a brief overview of the results that have been achieved in this rapidly developing field. We start by describing how human players deviate from conventional rationality. Three aspects can be distinguished: (1) motivation (what drives people; how do players evaluate outcomes?); (2) cognition (how do people reason; what thinking processes do they use when they are confronted with a new game?); and (3) adaptation (how do people learn when they play the same game repeatedly?). After having described empirical regularities, we briefly discuss novel theories relating to each of these aspects.

Bounds on Human Rationality

Humans have emotions, face cognitive limitations, and may be under (time) pressure; they process information differently, or may be motivated differently than standard theory assumes. Conventional theory is silent on the decision-making processes

that are used. It assumes that players only differ in preferences, not in cognitive abilities. Models of bounded rationality take into account limits on human knowledge, informational processing capacity and computational ability. Obviously, in some games, these limits are more important than in others.

When observed behavior differs from the game theoretic prediction, one can point to one of two main causes:

1. The game, as perceived by the players, is different from the one analyzed by the theorist.
2. The solution concept is not applicable in this situation.

Conventional game theory starts with the model and assumes that it is common knowledge among the players. Real life conflict situations are less structured and have to be interpreted; a model has to be constructed, and different players may perceive the situation differently. At first, one does not necessarily see all aspects of the problem, leading to superficial decision making. Nevertheless, this may already produce a satisfactory solution, not inviting further reflection. [Selten \(1998\)](#) notes that such superficiality may explain the framing effect ([Tversky and Kahneman, 1981](#)): the way the situation is presented may have an important influence on the outcome. The reasoning process anchors at aspects that deeper inspection might reveal to be irrelevant. Alternatively, framing may provide clues to the solution that conventional theory mistakenly neglects; [Schelling \(1960\)](#).

Traditional models assume that players are rational in the sense of Von Neumann and Morgenstern (expected utility) and Savage (subjective expected utility). Mostly, the assumption of players having a common prior is added. Following the path-breaking [Kahneman and Tversky \(1979\)](#), behavioral economics has shown that humans deviate from these assumptions and the resulting behavior in systematic ways: (1) utility depends not just on the final state (outcome), but on changes in the state; (2) the change is measured with respect to a reference point; (3) losses loom larger than gains (loss aversion); (4) individuals use decision weights that are different from probabilities, with small probabilities being overweighted and large ones underweighted; and (5) information is processed differently than Bayes' rule describes (see the article on Behavioral Economics for more details). Furthermore, while the rational choice model allows general preferences, in empirical work it is frequently assumed that players are selfish and care only about own materialistic payoffs. Many humans are motivated differently: altruism and reciprocity play a role, with (social) norms influencing behavior as well.

Equilibrium concepts assume that (1) players form beliefs about what others will do, (2) these beliefs are correct, and (3) players best respond to them. In equilibrium, players never are surprised. A game may be too complicated to find a best response, or players may be insufficiently motivated to find one. Equilibrium concepts are based on circular reasoning (fixed points; solutions to systems of equations), but, as stressed by [Selten \(1998\)](#), humans have a tendency to avoid circular concepts. The natural way of problem solving is by using step-by-step reasoning processes. The rationalistic interpretation of equilibrium assumes, but leaves unexplained how the beliefs of players that are confronted with a new game come to be correct, an assumption that is

especially problematic in games with multiple equilibria. Finally, natural learning processes need not converge to (Nash) equilibrium, or the learning may be too slow to be practically relevant.

Although for certain classes of games, nonequilibrium concepts, such as rationalizability or iterated dominance, are sufficient to predict the outcome, these rely on an unlimited number of iterations; humans seem to do only a very small number of rounds of iterated strategic thinking. Related, humans find processes like backward induction to be unnatural and do not always use these.

Empirical Regularities

Bounded rationality (the rationality displayed by humans in decision-making situations), hence, differs significantly from perfect rationality. As a result, it is not surprising that experiments have revealed that outcomes observed when humans play games differ systematically from the standard game theory predictions; see [Camerer \(2003\)](#), [Goeree and Holt \(2001\)](#), and [Selten \(1998\)](#). The following are some observed empirical regularities:

1. Framing effects can be very important, even in simple zero-sum games.
2. The outcome may depend on aspects of the game that conventional theory considers to be irrelevant; in contrast, aspects that the theory considers relevant need not matter; for example, two games with the same unique mixed strategy equilibrium may be played differently.
3. Not only the ordering of the payoffs matters, but also payoff differences: strategies that are not best responses may be played and small payoff differences may be ignored altogether.
4. Frequently, players care about other aspects than own (material) payoffs.
5. In games played once, the observed outcome may differ from the unique rationalizable one; in strategic games, people only do a small number of rounds of iterated elimination.
6. In perfect information games, monitoring of the decision making process shows that players may not do backward induction; this procedure does not come naturally, but it can be taught.
7. When players gain experience with a game, they adjust behavior; players learn in different ways and with different speeds, which may depend on the game; learning processes may be very slow.

Players' Motivations

The ultimatum game ([Güth et al., 1982](#)), along with its variants, the dictator game and the trust game, has spawned a large literature on social preferences. In the ultimatum game, P_1 proposes how to divide an amount of money between him and P_2 ; if P_2 accepts, the proposal is implemented; otherwise, each player gets 0. If both players are selfish (care only about own material payoffs), P_1 offers (close to 0) to P_2 , which the latter accepts. Experiments, however, show that most offers are close to 50/50 and that the responding player may reject

considerable amounts. The dictator game (the variant in which P_2 is forced to accept all proposals) allows to test whether the behavior of P_1 is driven by altruism or by the fear that low offers will be rejected. In the trust game, P_1 can first transfer all or part of the amount to P_2 , with that amount, T , being multiplied by a known constant $K > 1$, and with P_2 subsequently deciding how much to transfer back. In this game, responding players display positive reciprocity: the larger the transfer received, the more is transferred back. In public goods games, we see negative reciprocity: players that do not contribute sufficiently to the public good are punished. All these games show considerable individual heterogeneity, while culture matters as well; [Henrich et al. \(2004\)](#).

While the experimental results do not refute game theoretic analysis as such, they suggest that great care is needed in modeling a situation as a game; the frequently made assumption of selfishness does not describe most situations very well. Player 2 may be motivated to get a reasonable share and he may prefer conflict to an outcome in which he gets much less than the proposer. The game of [Figure 1](#) is a kind of mini-ultimatum game: P_2 may reject if Y is too small. On the other hand, if $X \cong 0$ and $Y = 1$, P_2 might realize that P_1 is forced to divide asymmetrically; he may accept in this case, while he might possibly reject if X is large and $Y \cong 0$. Hence, not only the outcomes matter, but also the context in which these arise; consequentialism is violated. In the literature, a large variety of models has been proposed and tested to incorporate these aspects, including models of pure distributional preferences, as well as models in which intent or procedural aspects matter. Given heterogeneity, incomplete information about preferences is natural, and there are also models in which players care about what others think of them.

[Sobel \(2005\)](#) is a recent overview of the literature on interdependent preferences.

In many models that fit under this heading, just the players' payoff functions are changed by incorporating behavioral aspects, while the solution concept remains conventional. [Geanakoplos et al. \(1989\)](#), however, argued that emotions cannot be captured by assuming that payoffs just depend on players' actions. Players' emotions, which influence how outcomes are valued (payoffs), will also depend on their expectations and, hence, on what they learn in the game. For example, in [Figure 1](#), if P_2 expects P_1 to choose O , then he may be disappointed when asked to move and, therefore, choose R , while a P_2 in that expected I may simply play L . [Geanakoplos et al. \(1989\)](#) introduce 'psychological games' as games in which payoffs depend both on what players do and on what they think and they define a 'psychological equilibrium' as a profile in which each player best responds and beliefs are correct. [Rabin \(1993\)](#) builds on this idea to construct a model that incorporates considerations of fairness. Starting from an ordinary game and players' expectations, he first applies a kindness function to modify the payoffs to take into account the emotions and construct a psychological game, to which then the concept of psychological equilibrium is applied. The resulting outcome is called a fairness equilibrium. In the game from [Figure 1](#), both the outcome (1, X) and the outcome (2, Y) can be supported by fairness equilibria.

Modeling Players' Thinking Processes

The p -beauty contest game ($0 < p < 1$) illustrates the difference between individual human rationality and rationalizability. In this game, n players are asked to simultaneously pick an integer from $\{1, 2, \dots, 100\}$, with the person whose number is closest to p times the average being the winner (ties broken randomly). The unique Nash equilibrium is for all to choose 1; in fact, if $p < 3/4$, only 1 is rationalizable. Experiments show that players that are new to the game exhibit distinct, bounded levels of reasoning. For example, for $p = 2/3$, games played with a large population of players, show spikes at numbers such as 50 (superficial thinking), 33 (the best response against 50), 22, 15, 10, etc. Obviously, not all players are perfectly rational, and in such situations, a fully rational player should not necessarily pick 0.

If one knows that other players are not that sophisticated and engage only in limited thinking, a natural idea is to try to figure out what naïve players might do and to estimate the distribution of 'levels of thinking' in the population, and to best respond against the resulting beliefs. For example, if one thinks that naïve, level-1 and level-2 players are equally likely, one believes the average chosen by the others to be 35 and chooses 23. Theories of cognitive hierarchy ([Camerer, 2003](#)) or level- k thinking have been developed, and have been brought to the data with reasonable success; see [Crawford et al. \(2013\)](#) for a recent overview.

Some insight has also been gained on the question of how 'naïve play' is obtained: players come to the laboratory with prior ideas and may seek for analogies with known situations, or may seek for clues, such as labels and focal points; [Schelling \(1960\)](#). It has also been shown that backward induction reasoning does not come natural; for example, eye tracking studies show that many players only look for a limited time, or not at all, at the parts of the game tree that come later; hence, there is a tendency for myopic decision making.

Learning

When players get more experience with a game, they will typically adjust their behavior. As already suggested in [Nash \(1950\)](#), learning may produce outcomes that are more in line with equilibrium predictions. For example, in repetitions of p -beauty contest games, smaller numbers are chosen over time. We already noted that there is a very large theoretical literature on learning. In broad terms, two types of learning models can be distinguished. In belief learning, players adjust their beliefs about what others will do and best respond to these. In reinforcement learning, a player's own strategies have certain 'attractions,' which determine the probability of playing them and which are adjusted through time. [Camerer \(2003\)](#) describes a hybrid model, experience-weighted attraction learning that has successfully been fitted to the data of several games. The data and model show that learning is much quicker in some games than in others, hence, that players come to certain games with rather strong, but not necessarily correct, beliefs. Of course, not all people learn in the same way and some people may learn quicker than others. [Selten \(1998\)](#) stresses that learning may involve a combination of adaptive and analytical approaches.

Players that view the situation as routine and use adaptive methods may fail to acquire an overview of the situation and, hence, may not be able to reach equilibrium; for that, a more analytical approach may be needed.

Conclusion

The allotted space only allows making just a few remarks on the applications of noncooperative game theory; the interested reader is suggested to consult the Handbook of Game Theory or to follow any of the leads given below.

One application is to cooperative game theory. The 'Folk Theorem' shows that when a game is repeated sufficiently often, players are sufficiently patient and obtain enough information about past play (and when some other technical conditions are satisfied), the set of subgame perfect equilibrium outcomes coincides with those of the static cooperative game. Hence, in a context of repeated interaction, contracts may not be necessary to achieve efficient outcomes. Work of Nobel Prize winner Elinor Ostrom on the exploitation of common property resources combines this theory with rich institutional data. Avner Greif's contribution to the Handbook surveys the related literature that employs game theory for economic history analysis.

Following Nash's (1950) suggestion, the literature has pursued the 'Nash program' and has shown that several cooperative solution concepts can be implemented non-cooperatively, i.e., there exists a noncooperative game of which the (refined) equilibrium outcomes correspond to those predicted by the cooperative solution concept. For example, the Nash-bargaining solution can not only be implemented by means of the simultaneous bargaining game proposed already in Nash (1953), but also by a very natural alternating offers procedure proposed in Rubinstein (1982). In a similar vein, solution concepts such as the Shapley value, the kernel, and the nucleolus can be obtained by noncooperative procedures; we refer to Thomson's article in this Encyclopedia for details.

Within economics, the field of industrial organization has proved a fertile ground for application. Probably the most successful applications thus far have been to auctions or more generally market design; see the article on Auctions. Following the seminal work of Vickrey, auction theory first focused on the questions of how to bid in a single-item auction and what is the revenue maximizing mechanism for selling such item. In implementing spectrum policy, governments were, however, confronted with the question of how to sell multiple heterogeneous items, given certain objectives, such as market efficiency and raising sufficient government revenue. Research, in which game theorists interacted with experimental economists and Operations Researchers, has led to highly innovative auction designs, such as the simultaneous multi-round auction and the combinatorial clock auction, which have subsequently been successfully implemented by various governments around the world.

By using both cooperative and noncooperative approaches, and by combining theory with experimental and empirical studies, Nobel Prize winner Alvin Roth and coworkers showed that stability is important for

understanding the success of particular market institutions. Building on this insight, they successfully reengineered several existing institutions, such as those for matching organ donors with patients. Importantly, this work shows that economics can take into account ethical restrictions, such as the prohibition of side payments.

As seen above, game theory has provided a strong stimulus for experimental economics. No doubt, the further development of behavioral game theory will give a further boost to applications.

See also: Auctions; Behavioral Economics; Cooperative Game Theory; Experimental Economics; Information, Economics of; Old and New Institutionalism in Economics; Transaction Costs and Property Rights.

Bibliography

- Aumann, R.J., 1987. Game theory. In: Eatwell, J., Milgate, M., Newman, P. (Eds.), *The New Palgrave Dictionary of Economics*, vol. 2, pp. 460–482.
- Aumann, R.J., 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 67–96.
- Aumann, R.J., 1976. Agreeing to disagree. *Annals of Statistics* 4, 1236–1239.
- Aumann, R.J., Hart, S., 1992/1994/2002. *Handbook of Game Theory with Economic Applications*, vol. 1, pp. 1–731, vol. 2, pp. 735–1520, vol. 3, pp. 1521–2351, Elsevier, North Holland.
- Bernheim, B.D., 1984. Rationalizable strategic behavior. *Econometrica* 52, 1007–1029.
- Camerer, C., 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ.
- Carlsson, H., Van Damme, E., 1993. Global games and equilibrium selection. *Econometrica* 61, 989–1018.
- Crawford, V.P., Costa-Gomes, M.A., Iriberrí, N., 2013. Structural models of nonequilibrium strategic thinking: theory, evidence, and applications. *Journal of Economic Literature* 51 (1), 5–62.
- Fudenberg, D., Levine, D., 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- Fudenberg, D., Levine, D.K., 2009. Learning and equilibrium. *The Annual Review of Economics* 1, 385–419.
- Geanakoplos, J., Pearce, D., Stacchetti, E., 1989. Psychological games and sequential rationality. *Games and Economic Behavior* 1 (1), 60–79.
- Goeree, J.K., Holt, C.A., 2001. Ten little treasures of game theory and ten intuitive contradictions. *The American Economic Review* 91 (5), 1402–1422.
- Govindan, S., Wilson, R., 2008. Metastable equilibria. *Mathematics of Operations Research* 33 (4), 787–820.
- Güth, W., Schmittberger, R., Schwartz, B., 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3 (4), 367–388.
- Harsanyi, J., 1967–1968. Games with incomplete information played by Bayesian players, parts I, II and III. *Management Science* 14, 159–182, 320–334, 486–502.
- Harsanyi, J., 1973. Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory* 2, 1–23.
- Harsanyi, J., Selten, R., 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-scale Societies*. Oxford University Press, Oxford, UK.
- Kagel, J., Roth, A., 1995. *Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ.
- Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decisions under risk. *Econometrica* 47 (2), 263–291.
- Kohlberg, E., Mertens, J.F., 1986. On the strategic stability of equilibria. *Econometrica* 54, 1003–1039.
- Kreps, D., Wilson, R., 1982. Sequential equilibria. *Econometrica* 50, 863–894.

- Kuhn, H., 1953. Extensive games and the problem of information. In: Kuhn, H., Tucker, A.W. (Eds.), *Contributions to the Theory of Games II*. Princeton University Press, Princeton, NJ, pp. 193–216.
- Mertens, J.F., 1989. Stable equilibrium: a reformation, part I: definition and basic properties. *Mathematics of Operations Research* 14, 575–625.
- Morris, S., Shin, H., 2003. Global games: theory and applications. In: Dewatripont, M., Hansen, L., Turnovsky, S. (Eds.), *Advances in Economics and Econometrics* (Proceedings of the Eighth World Congress of the Econometric Society). Cambridge University Press, Cambridge, UK.
- Myerson, R.B., 1991. *Game Theory*. Harvard University Press, Cambridge, MA.
- Nash, J., 1950. Non-cooperative Games. Ph.D. dissertation. Princeton University, Princeton, NJ.
- Nash, J., 1951. Non-cooperative games. *Annals of Mathematics* 54, 298–395.
- Nash, J., 1953. Two-person cooperative games. *Econometrica* 21, 128–140.
- Osborne, M.J., Rubinstein, A., 1994. *A Course in Game Theory*. MIT Press, Cambridge, MA.
- Pearce, D., 1984. Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52, 1029–1051.
- Rabin, M., 1993. Incorporating fairness into game theory and economics. *American Economic Review* 83 (5), 1281–1302.
- Rubinstein, A., 1982. Perfect equilibrium in a bargaining model. *Econometrica* 47, 1353–1366.
- Rubinstein, A., 1989. The electronic mail game: strategic behavior under 'almost common knowledge'. *American Economic Review* 79, 385–391.
- Schelling, T., 1960. *The Strategy of Conflict*. Harvard University Press, Cambridge, MA.
- Selten, R., 1965. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragertragheit. *Zeitschrift für die gesamte Staatswissenschaft* 121, 301–324, 667–689.
- Selten, R., 1975. Re-examination of the perfectness concept for extensive form games. *International Journal of Game Theory* 4, 25–55, 19.
- Selten, R., 1998. Features of experimentally observed bounded rationality. *European Economic Review* 42, 413–436.
- Smith, J., 1982. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK.
- Sobel, J., 2005. Interdependent preferences and reciprocity. *Journal of Economic Literature* 43 (2), 392–436.
- Tversky, A., Kahneman, D., 1981. The framing of decisions and the psychology of choice. *Science* 211 (4481), 453–458.
- Von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*, third ed. Princeton University Press, Princeton, NJ.
- Weibull, J., 1995. *Evolutionary Game Theory*. MIT Press, Cambridge, MA.
- Young, H.P., 2004. *Strategic Learning and Its Limits*. Oxford University Press, Oxford, UK.
- Zermelo, E., 1913. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiel. *Proceedings of the 5th International Congress of Mathematicians* 2, 501–504.