



Game theory

Martin Dufwenberg*

Game theory is a toolkit for examining situations where decision makers influence each other. I discuss the nature of game-theoretic analysis, the history of game theory, why game theory is useful for understanding human psychology, and why game theory has played a key role in the recent explosion of interest in the field of behavioral economics. © 2010 John Wiley & Sons, Ltd. *WIREs Cogn Sci*

INTRODUCTION

Game theory is a toolkit for analyzing situations where decision makers influence each other. Parlor games such as chess, bridge, and poker can exemplify this, and this relation explains the subject's name. However, game theory's main applications concern society: economists study, for example, auctions, bargaining, market competition, or family formation; military strategists examine conflict; biologists model natural selection; political scientists analyze voting, etc.

Game theory offers ways to compactly describe such situations. The resulting games are often fascinating, highlighting key aspects of interesting situations. Economists, psychologists, sociologists, and neuroscientists increasingly use such games to learn about human nature through experiments. This helps evaluate theories of, for example, bounded rationality, various emotions, or brain activity, or may inspire new theories. Such insights square well with game theory as the framework is flexible enough to allow many ways to model various psychological considerations.

In what follows, I will attempt to elucidate these aspects of game theory. My method of presentation will include several examples, in the hope that this will make the presentation vivid. Toward the end, I give a brief history of game theory, with particular attention to a historical fluke, which I believe gives game theory a key role in explaining the recent explosion of interest in the field of behavioral economics.

EXAMPLES OF GAMES

Example 1 (*the game of 21*) Two persons, call them A and B, take turns making choices. A begins. To

start off, he can choose either 1 or 2. B observes this choice, then increments the 'count' by adding one or two. That is, if A chooses 1 then B can follow up with 2 or 3; if A chooses 2 then B can follow up with 3 or 4. A then observes B's choice, and again increments the count by adding one or two. The game continues with the players taking turns, incrementing the count by one or two. The player who reaches 21 wins.

The game of 21 might be played circling numbers on the following strip:

| | | | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 |
|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|

Imagine that you are player B. What would you do? Think about it for a minute before reading on. The exercise may help you appreciate more deeply some of the comments to follow.

A game-theoretic analysis might proceed as follows: First, identify the decision makers in the game, called players. Here they are A and B. Second, identify the strategies for each player, meaning comprehensive descriptions of what a player would do at every conceivable situation that might arise in the game. For example, a strategy for player B states which choices he would make at each of the counts 1, 2, 3, ..., 20, conditional on any sequence of choices up to that count. Finally, add a description of how strategies combine to produce outcomes and specify the players' preferences over these outcomes. In this case, while there are a great number of strategy combinations each of these results in one of only two outcomes, specifying whether A or B wins. A natural assumption about preferences is that each player prefers to win.

The number of strategies a player has in the game of 21, as defined in the previous paragraph, is bewildering. Player B, for example, has (if I calculated correctly) 2^{8855} of them. This is far more than

*Correspondence to: martind@eller.arizona.edu

Department of Economics, University of Arizona, Tucson, AZ 85750-0108, USA

DOI: 10.1002/wcs.119

the estimated number of atoms in the universe. Is there any hope of generating a meaningful prediction for the game? The answer is *yes*. Game-theorists have terminology for describing details regarding the order in which players move as well as the information they have when they do so. The game of 21 is then a finite, two-player, zero-sum, perfect information, two-outcome game with perfect recall. This is not the place to formally define these terms, but it follows from the classification that one of the players can force a win regardless of the other's choices.¹ Specifically, B can guarantee a win by always choosing a multiple-of-three, thus securing that he will make the sequence of choices 3, 6, ..., 21. A game theorist would say that this is part of a *dominant strategy*, a strategy with the property of being no worse than any other strategy regardless of the strategy choice of the opponent. Not all games admit dominant strategies, but the game of 21 does (for player B).

The remarks of the previous paragraphs give some flavor of how game-theoretic analysis may proceed. It is not obvious, however, that the described solution is an empirically accurate reflection of how humans play (honestly, what was your choice?). Indeed, experiments reveal that most humans often fail to choose a dominant strategy in this or similar games, at least until they get some practice (see Refs 2,3; cf. also Ref 4 which ran experiments on Nim, an old game with some similar features first analyzed theoretically more than a century ago by Bouton⁵). Through such experiments one may gain insight about human cognition and the psychology of learning.

Two-player games where players have few strategies (much fewer than 2⁸⁸⁵⁵) are often described using game matrices. Here is an example where each player has exactly two strategies:

Example 2 (stag-hunt)

| | | |
|-------------|-------------|-------------|
| | <i>Stag</i> | <i>Hare</i> |
| <i>Stag</i> | 9, 9 | 0, 8 |
| <i>Hare</i> | 8, 0 | 7, 7 |

Two players A and B simultaneously make choices *Stag* or *Hare*. A's strategies correspond to rows while B's strategies corresponds to columns. The numbers in brackets indicate the players' payoffs as a function of their choices, with player A's payoff stated first. For example, if A chooses *Stag* while B chooses *Hare* then A gets 0 and B gets 8.

The stag-hunt game got its name from a story involving two hunters who independently choose

between hunting stag (a risky choice which pays off only if the other hunter goes stag-hunting too) or hare (a safer choice). If the numbers in the matrix were different we could get different two-player two-strategy games and many of these have different and colorful names. The following games are, respectively, versions of games known as *prisoners' dilemma*, *chicken*, *battle-of-the-sexes*, and *matching pennies*:

Examples 3-6

| | | | | | |
|----------|----------|----------|--|----------|----------|
| | <i>C</i> | <i>D</i> | | <i>X</i> | <i>Y</i> |
| <i>C</i> | 2, 2 | 0, 3 | | 0, 0 | 4, 1 |
| <i>D</i> | 3, 0 | 1, 1 | | 1, 4 | 3, 3 |
| | <i>B</i> | <i>F</i> | | <i>H</i> | <i>T</i> |
| <i>B</i> | 3, 1 | 0, 0 | | 1, -1 | -1, 1 |
| <i>F</i> | 0, 0 | 1, 3 | | -1, 1 | 1, -1 |

Some of these games are frequently taken to provide useful metaphors for societal analysis. For example, in the prisoners' dilemma *D* is a dominant strategy for each player. But if each player chooses *D* then they both end up worse off than had each chosen *C* (they get 1 instead of 2). This illustrates a potential conflict between individual rationality and the common good, a theme which may have important societal implications more broadly. Various scholars have argued that situations involving, for example, pollution, arms races, or advertising often give rise to prisoners' dilemma-like situations. Or see Myerson⁶ who uses stag-hunt games to discuss 'cultural roots of poverty' and chicken games to discuss 'property rights and justice'; I would be remiss if I did not here refer also to Thomas Schelling's classic *The Strategy of Conflict*,⁷ from which Myerson's says he is 'learning'.

NASH EQUILIBRIUM

While a game is a description of a strategic situation, a solution concept is a mathematical rule for generating predictions. The notion of a dominant strategy, mentioned above, can exemplify (the prediction being that a player would choose a dominant strategy if he has one). However, not all games possess dominant strategies. For example, neither player has a dominant strategy in the stag-hunt, chicken, battle-of-the-sexes, or matching pennies games.

The most important and well-known solution concept in game theory is the *Nash equilibrium*: a combination of strategies, one for each player, with the property that each player's strategy is optimal given each other player's choice. Example 2,


for example, has three Nash equilibria: each player chooses *Stag*; each player chooses *Hare*, and finally an equilibrium where each player uses a *mixed strategy* (which is a probability distribution over pure strategies) where *Stag* is chosen with probability $7/8$ and *Hare* with probability $1/8$. Note that a Nash equilibrium describes the behavior of all players in a game, unlike the notion of a dominant strategy which relates to an individual player.

Much game-theoretic analysis, not the least applications in economics, concerns Nash equilibria. Since a game may have multiple Nash equilibria, it is natural to wonder which one is most plausible. The literature on equilibrium refinement/selection^{8–10} develops criteria, based on armchair reasoning, to obtain more specific predictions. In the stag-hunt game, for example, it is natural to wonder whether players will coordinate on the higher-payoff equilibrium (*Stag, Stag*) or the lower-payoff equilibrium (*Hare, Hare*) which from each player's viewpoint may seem less risky. Different theories provide different answers. And what if we augment the game so that players can talk to each other? Aumann¹¹ argues that this might not help the players coordinate on the good stag-hunting equilibrium. The reason would be that even a player planning to choose *Hare* would want to convince the other player to choose *Stag*, since he would therefore increase his payoff from 7 to 8. Therefore any claim that a player intends to play *Stag* may not be credible, so that communication would not help. The empirical relevance of the involved ideas here—which equilibrium is most plausible in the first place and whether communication helps—has been tested in experiments.^{12,13}

THE EXTENSIVE FORM AND BACKWARD INDUCTION

Within game theory there are different models. One specification—implicit in Examples 2–6—refers only to players, strategies, and preferences and is called the *normal form*. In many games, including the game in Example 1 above, players move in turns and maybe many times. While the normal form is not explicit about such matter, the alternative model of games represented in *extensive form* (originally proposed by von Neumann and Morgenstern¹⁴ and further developed by Kuhn¹⁵) provides such information and allows for quite intuitive depictions via game trees. Here is an example:

Example 7 (*hold-up game*)



This game comes with a story¹⁶: An artist (A) has been asked by a presumptive buyer (B) to paint a ‘beautiful portrait of B’. A may say yes or no to accepting to do the job. In the former case, A and B go separate ways. In the latter case, A spends \$2000 worth of her time on the painting, and a contract says B should subsequently pay \$5000 to A. The value to B is \$8000 but B may complain and claim (falsely) that the portrait is ‘rather ugly’ and attempt to renegotiate offering a new price of \$1000. Given the unclear definition of beauty, A cannot enforce the \$5000 payment and will have to accept or reject the new offer. A knows that no person other than B would pay to get the painting.

The game tree captures this story. Play starts at the top, with branches *No* and *Yes* reflecting A’s initial choice, etc. The terminal nodes, where numbers appear, indicate that play ends and the numbers reflect the players’ monetary payoffs as reflected in the story (divided by 1000), with player A’s payoff on top. For example, if A chooses *Yes*, if B then attempts to renegotiate (= chooses *Reneg*), and if A then chooses *Accept*, then A’s payoff is -1 $[= (1000 - 2000)/1000]$ and B’s payoff is 7 $[= (8000 - 1000)/1000]$, etc.

An economist inclined to assume that people maximize their monetary earnings would solve the game of Example 7 using the game-theoretic notion of backward induction: Following choices *Yes* and *Reneg*, A would choose *Accept* since $-1 > -2$. Knowing this, following choice *Yes*, B would choose *Reneg* since $7 > 3$. Knowing this, at the start of the game, A would choose *No* since $0 > -1$. The outcome is that each gets 0; too bad for them since each could have had 3 had they chosen *Yes* and *Pay*.

ENRICHING THE PSYCHOLOGY

Example 7 illustrates what a game tree is and the notion of backward induction. It also illustrates (in light of the story) the economic insight that partnership-specific investments and incomplete contracts may combine to cause an inefficient outcome; this is what economists call hold-up. In addition, it can illustrate how games may inspire insights of how economic analysis may be fruitfully enriched in psychological directions. In particular, keeping in mind the artist–buyer story, it seems plausible that once the buyer renegotiates the artist would view A as unkind and therefore perhaps choose *Reject* in order to hurt B and get even. If B foresees this he may therefore choose *Pay*; if A foresees that he may choose *Yes*. Take into account the desire for A to get even and we can conclude that *every choice may flip*. A vengeance motive may thus help bring about a rosy outcome with realized gains-from-trade.

Economists used to working with formal theory would want to recast the intuitive analysis of the previous paragraph using a formal model. This has been done¹⁶ by applying a model of *reciprocity*,^{17,18} which furnishes a way to capture formally the idea that people desire to reward kindness with kindness and to avenge unkindness. This is not the place to describe the details of this theory, but I note that it builds on the toolbox of game theory. While game theory was not originally developed in order to model reciprocity, it turns out that game theory provides a flexible enough platform that one can proceed in that direction.

This last point can be generalized. Pick almost any exciting psychological direction (say regarding bounds of humans reasoning capabilities, psychological biases, or complex forms of motivation) in which one may wish to enrich standard economic analysis (which typically builds on attributing to decision makers a desire to maximize own income as well as unbounded reasoning capabilities). I propose that game theory is likely to provide a flexible enough platform that one can address the consideration analytically.

Let me run through a couple of examples to support that claim: First, several game-theory-based models have recently been proposed where people care not only about own income but rather about the overall distribution of income. Fehr and Schmidt¹⁹ discuss the approach, which has also been applied to hold-up scenarios like that in Example 7.²⁰ Second, a broader framework, introduced and baptized *psychological game theory* by Geanakoplos et al.²¹ and further developed by Battigalli and Dufwenberg²² provides the intellectual home of reciprocity theory (of

the kind mentioned above) as well as a way to model a variety of emotions (including disappointment, guilt, anger, regret, frustration, and anxiety) and also sociological concerns like esteem and social respect. The framework differs from traditional game theory in that payoffs are directly influenced by beliefs (about strategies and beliefs), and not only by the strategies chosen. (The term ‘theory of games with belief-dependent motivation’ would be more descriptive than the term ‘psychological game theory’, but the latter has become established, so I use it here.) Psychological game theory is thus not a special case of game theory, but rather an extension. Third, fourth, and fifth models concerning learning (how players change their strategic choices with experience), cognitive hierarchies (how players reason about one another), and bounded rationality (how players fail to optimize) have been explored starting from game-theoretic platforms. Space constraints prevent me from going into details but for more on these three respective topics, see chapters 6 and 5 of the work of Camerer²³ and Rubinstein²⁴.

The range of examples here is limited. Yet I feel the list is long enough to justify the general belief that game theory-based approaches to psychology constitute a potentially fruitful area for future research.

THE HISTORY OF GAME THEORY

This entry has by now morphed to focus on the relevance of game theory to human psychology. That is not how game theory started, and I now move to give a brief account of the history of game theory. This discussion will eventually (in the next section) lead back to a cognitive science theme.

Modern game theory starts in 1928 with John von Neumann’s²⁵ analysis of two-person zero-sum games, where one player’s gain equals the other’s loss (as in poker, chess, the game of 21, matching pennies, and possibly war). He proposed solutions with some remarkable properties. However, in collaboration with Oskar Morgenstern he realized that many interesting economic situations lead to games which have more players or are not zero-sum. Examples 2–5, and 7 exemplify. In the prisoners’ dilemma, for example, both players are better off when both play C than when both play D, so the game is not zero-sum. In their book *The Theory of Games & Economic Behavior*, von Neumann and Morgenstern¹⁴ proposed methods for analyzing games generally and their pioneering contribution is probably the most important milestone in the history of game theory.

Following this work, game theory can be classified into two main strands. The first is *non-cooperative game theory* and comprises mainly games represented in normal form as well as extensive form. This approach is exemplified by Examples 1–7. The other approach is *cooperative game theory*, which uses models in which strategic details are not given. Rather, focus is on coalitions (= subsets) of players and what they can achieve, and a common interpretation is that cooperative game theory studies the outcome of joint actions in a situation with external commitment (which would be absent unless formally modeled in the non-cooperative approach). Von Neumann and Morgenstern's approach is non-cooperative as regards two-player zero-sum games and cooperative for other games. A large literature on cooperative games have subsequently developed. One example of an important contribution is the so-called Nash bargaining solution. Rather than specify a non-cooperative game with explicit strategies, Nash²⁶ postulates a few axioms regarding properties that binding agreements may be expected to possess (for example, if the agreement entails some distribution of payoffs x then there cannot be another agreement with payoffs y such that each bargaining party prefers y to x). Those axioms turn out to pin down a particular distribution of payoffs, which is the proposed solution.

From now on I will not discuss cooperative game theory any more and instead refer interested readers to Aumann's²⁷ entry on game theory in the *Palgrave: Dictionary of Economics*, which has a historical structure and covers a lot of the developments. I move to another milestone in the history of game theory: John Nash's article *Non-Cooperative Games*,²⁸ which defines and proves the existence of a Nash equilibrium in pure or mixed strategies in any game with finitely many players and strategies. The main reason why Nash's article is seminal is, however, perhaps not this mathematical result, but rather the underlying motivation for his study that reflected an understanding that it was fruitful to study all games, not only zero-sum games, in a non-cooperative mode and that Nash equilibria would be useful in this connection. Judged by the frequency with which the concept has been applied in economics as well as in other fields, the notion of Nash equilibrium is clearly the most important game-theoretic notion around.

Nash received the Nobel Memorial Prize in Economics in 1995, together with John Harsanyi and Reinhard Selten, for 'for their pioneering analysis of equilibria in the theory of non-cooperative games'. Selten was celebrated for his work on equilibrium refinements (*op. cit.*), and Harsanyi for his work²⁹ extending the Nash equilibrium notion to games

with incomplete information (where certain game features may not be known to everyone; this is important for example in auctions where a bidder may not know another's evaluation). See Ref 30 for more discussion of these contributions. Several other Nobel Memorial Prizes have since gone to scholars who were celebrated for their game theory related contributions: Robert Aumann and Thomas Schelling shared the Prize in 2005 'for having enhanced our understanding of conflict and cooperation through game-theory analysis'. Leonid Hurwicz, Eric Maskin, and Roger Myerson shared the Prize in 2007 'for having laid the foundations of mechanism design theory', a field which draws heavily on game theory. Moreover, the Prizes in 1996 to James Mirrlees and William Vickrey and in 2001 to George Akerlof, Michael Spence, and Joseph Stiglitz concerned the economics of 'asymmetric information', and much of the celebrated work draws heavily on game theory. For more details about all of this work, I refer to the many well-written texts available on the Nobel Foundations' site: http://nobelprize.org/nobel_prizes/economics/

One important set of game-theoretic models that deserves special mention, used and originally developed by biologists to study natural selection, is *evolutionary game theory*. The basic idea is this: When organisms play some game of life each organism is hard-wired to use a certain strategy. The resulting payoffs represents fitness, or number of offspring, so more successful behaviors increase their share of the population over the generations. Two classic references are Maynard-Smith and Price³¹ and Maynard-Smith.³² Subsequently it has been argued that models of this sort can be relevant to economics (for example, viewing markets as games, firms as the interacting organisms, and profit as fitness) and economists have done much of the further theory development. See Ref 33 for a synthesis and overview of much of the literature.

BEHAVIORAL ECONOMICS & GAME THEORY

What is the most important more recent development related to game theory? Aumann's²⁷ dictionary entry on game theory (written for an earlier edition) stops at 1986. *What has happened since?* Different scholars would surely stress different matters. To describe my view, I must first let go of my focus on game theory and instead discuss another field.

Behavioral economics is concerned with incorporating into economics a psychologically richer conception of man than has been customary in classical work. Interest in behavioral economics has exploded

in the last decade. I cannot review all of this literature here; the references given in the section 'Enriching the Psychology' above provide a handful of examples. The anthology Camerer et al.³⁴ contains many others. Behavioral economics is not a new field. Arguments that conventional economic models do not provide a rich enough conception of man go way back and include contributions by scholars including Maurice Allais, Daniel Ellsberg, Herbert Simon, Daniel Kahneman, and Amos Tversky. Much of this work was celebrated and received attention, but nevertheless I think it is fair to say that for a long time the contributions did not really affect mainstream economic theory. This has changed dramatically the last 15 years. Here is the 10 thousand dollar question: *why?*

My answer puts game theory at center stage. The 1960s and early 1970s was a slow period for non-cooperative game theory, or at least it was a period where most economists did not pay much attention to game theory and were often not trained in game theory. In the mid-1970s and increasingly onwards this started to change, however. For example, the game theory based study of economic situations with asymmetric information (for example when employers have less information than job applicants about their talent, as in Spence's³⁵ job-market signaling model) took off in the 1970s. Bargaining theory, which is based on game theory, became popular, in particular through the contribution of Rubinstein.³⁶ The methodological core of the field of industrial organization changed

completely in nature and became game theory based; the textbook by Tirole³⁷ reflects that development well. New exciting results were derived on epistemic game theory, for example, concerning the implications common belief of rationality without assuming that players will coordinate on an equilibrium,^{38–41} on repeated games (where players interact over and over again in some given game⁴²), and on the impact of communication in games.^{43–45} Concomitant with these important contributions, interest in game theory grew. *Students were taught game theory.*

This happened because people were excited about asymmetric information, about bargaining, about industrial organization, about foundational topics in game theory, etc. This did not originally happen because people wanted to enrich economics psychologically and to be able to make progress in behavioral economics. However, as I have explained above, knowing game theory is a good start for anyone who wishes to incorporate a richer conception of man in economic analysis and to do exercises in behavioral economics. So when it happened that a new younger generation of economist was brought up on game theory, at some point many people in this generation realized that they were well positioned to explore topics in behavioral economics. Eventually this led to a boom in experimental and theoretical research addressing psychological aspects of game play and, more generally, to a boost of interest in behavioral economics. In my view this reflects the most important recent game theory related development.

ACKNOWLEDGMENTS

I thank Omar Al-Ubaydli, Geir Asheim, Pierpaolo Battigalli, Alan Sanfey, Alec Smith, and a referee for very helpful comments, and the NSF for financial support.

REFERENCES

1. Ewerhart C. Chess-like games are dominance solvable in at most two steps. *Games Econ Behav* 2000, 33:41–47.
2. Dufwenberg M, Sundaram R, Butler D. Epiphany in the Game of 21. *J Econ Behav Organ* 2010, 75:132–143.
3. Gneezy U, Rustichini A, Vostroknutov A. Experience and insight in the race game. *J Econ Behav Organ* 2010, 75: 144–155.
4. McKinney N, Van Huyck J. Does seeing deeper into a game increase ones' chances of winning? *Exper Econ* 2006, 9:297–303.
5. Bouton C. Nim, a game with a complete mathematical theory. *Ann Math* 1901–1902, 3:35–39.
6. Myerson R. Learning from Schelling's *Strategy of Conflict*. *J Econ Liter* 2009, 47:1109–1125.
7. Schelling T. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press; 1960.
8. Selten R. Reexamination of the perfectness concept for equilibrium points in extensive games. *Int J Game Theory* 1975, 4:25–55.
9. van Damme E. *Stability and Perfection of Nash Equilibria*. Berlin: Springer-Verlag; 1987.
10. Harsanyi J, Selten R. *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT-Press; 1988.
11. Aumann R. Nash equilibria are not self-enforcing. In: Gabszewicz J, Richard J-F, Wolsey L, eds. *Economic*

- Decision Making: Games, Econometrics & Optimization*. Amsterdam: Elsevier; 1990, 201–206.
12. Charness G. Self-serving cheap talk and credibility: a test of Aumann's conjecture. *Games Econ Behav* 2000, 33:177–194.
 13. Clark K, Key S, Sefton M. When are Nash equilibria self-enforcing? An experimental analysis. *Int J Game Theory* 2001, 29:495–515.
 14. von Neumann J, Morgenstern O. *Games and Economic Behavior*. Princeton: Princeton University Press; 1944.
 15. Kuhn H. Extensive games and the problem of information. *Ann Math Stud.* 1953, 28:193–216.
 16. Dufwenberg M, Smith A, Van M. Hold-up: with a vengeance. *Econ Inquiry* 2010, forthcoming.
 17. Rabin M. Incorporating fairness into game theory and economics. *Am Econ Rev* 1993, 83:1281–1302.
 18. Dufwenberg M, Kirchsteiger G. A theory of sequential reciprocity. *Games Econ Behav* 2004, 47:268–298.
 19. Fehr E, Schmidt K. Theories of fairness & reciprocity evidence & economic applications. In: Dewatripont M, Hansen LS, eds. *Turnovsky Advances in Economics & Econometrics: 8th World Congress of the Econometric Society*. Cambridge UP; 2003.
 20. Ellingsen T, Johannesson M. Promises, Threats, and Fairness. *Econ J* 2004, 114:397–420.
 21. Geanakoplos JD, Stacchetti E. Psychological games and sequential rationality. *Games Econ Behav* 1989, 1:60–79.
 22. Battigalli P, Dufwenberg M. Dynamic psychological games. *J Econ Theory* 2009, 144:1–35.
 23. Camerer C. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press; 2003.
 24. Rubinstein A. *Modeling Bounded Rationality*. Cambridge, MA: MIT Press; 1997.
 25. von Neumann J. Zur Theorie der Gesellschaftsspiele. *Math Ann* 1928, 100:295–293.
 26. Nash JF. The bargaining problem. *Econometrica* 1950, 18:155–62.
 27. Aumann R. Game Theory. In: Durlauf SN, Blume LE, eds. *The New Palgrave Dictionary of Economics*, vol. 3. 2nd ed. New York: Palgrave Macmillan; 2008, 529–558.
 28. Nash JF. Non-cooperative games. *Ann Math* 1951, 54:286–295.
 29. Harsanyi J. Games of incomplete information played by Bayesian players. Parts I, II, III. *Management Science* 1967–1968, 14:159–182, 320–334, 486–502.
 30. van Damme E, Weibull J. Equilibrium in strategic interaction: the contributions of John C. Harsanyi, John F. Nash and Reinhard Selten. *Scand J Econ* 1995, 97:15–40.
 31. Maynard-Smith J, Price GR. The logic of animal conflict. *Nature* 1973, 246:15–18.
 32. Maynard-Smith J. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press; 1982.
 33. Weibull J. *Evolutionary Game Theory*. Cambridge, MA: MIT Press; 1995.
 34. In: Camerer C, Loewenstein G, Rabin M, eds. *Advances in Behavioral Economics*. New York: Russell Sage Foundation; 2004.
 35. Spence M. Job market signaling. *Quart J Econ* 1973, 87:629–649.
 36. Rubinstein A. Perfect equilibrium in a bargaining model. *Econometrica* 1982, 50:97–110.
 37. Tirole J. *Industrial Organization*. Cambridge, MA: MIT Press; 1988.
 38. Pearce D. Rationalizable strategic behavior and the problem of perfection. *Econometrica* 1984, 52:1029–1050.
 39. Bernheim D. Rationalizable strategic behavior. *Econometrica* 1984, 52:1007–1028.
 40. Reny P. Common belief and the theory of games with perfect information. *J Econ Theory* 1993, 59:257–274.
 41. Battigalli P, Siniscalchi M. Strong belief and forward induction reasoning. *J Econ Theory* 2002, 106:356–391.
 42. Abreu D. On the theory of infinitely repeated games with discounting. *Econometrica* 1988, 56:383–396.
 43. Crawford V, Sobel J. Strategic information transmission. *Econometrica* 1982, 50:1431–1451.
 44. Forges F. An approach to communication equilibria. *Econometrica* 1986, 54:1375–1385.
 45. Myerson R. Multi-stage games with communication. *Econometrica* 1986, 54:323–358.
 46. Osborne M, Rubinstein A. *A Course in Game Theory*. Cambridge, MA: MIT Press; 1994.
 47. Watson J. *Strategy: An Introduction to Game Theory*. 2nd ed. New York: Norton; 2008.
 48. Dixit A, Skeath S, Reiley D. *Games of Strategy*. 3rd ed. New York: W. W. Norton; 2009.
 49. Dixit A, Nalebuff B. *The Art of Strategy: A Game-Theorist's Guide to Success in Business and Life*. New York: W. W. Norton; 2008.

FURTHER READING

There are many good books from which one can learn more about game theory. I use the textbook of Osborne and Rubinstein⁴⁶ when I teach on the PhD level, and that of Watson⁴⁷ when I teach upper-level undergraduates. If I were to teach a lower-level undergraduate class, I would use the book by Dixit et al.⁴⁸ Finally, for a lighter read about the art of strategy and a game theorist's guide to success in business and life, see Ref 49.