Difficulties in Analyzing Strategic Interaction: Quantifying Complexity

Chapter Seventeen

JUSTIN GRANA, RAND CORPORATION

This chapter is extracted from *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches*, by Aaron B. Frank and Elizabeth M. Bartels, eds., RR-A1275-1, 2022 (available at www.rand.org/t/RRA1275-1).

Prepared for the Defense Advanced Research Projects Agency,
Defense Sciences Office
Approved for public release; distribution unlimited

RAND
CORPORATION

For more information on this publication, visit **www.rand.org/t/RRA1275-1**.

Published by the RAND Corporation, Santa Monica, Calif.

CHAPTER SEVENTEEN

Difficulties in Analyzing Strategic Interaction: Quantifying Complexity

Justin Grana, RAND Corporation

Characterizing the complexity of strategic interaction remains an important challenge in the social sciences. One method for doing this is to apply the theory of computational-complexity and complexity classes to solving game-theoretic problems of strategic interaction. Although the mapping between strategic scenarios and computational-complexity classes is useful, it does not fully capture the dimensionality of strategic interaction that results from the presence or potential of multiple equilibria, communication, chaotic learning dynamics, and behavioral insights that affect how actors understand, analyze, and participate in strategic environments or games. In the pursuit of generality pertaining to computational needs, computational-complexity analysis may obfuscate the true difficulty in participating in or analyzing a strategic interaction. As a result, understanding the underlying structure of a particular game is crucial for refining the notion of complexity in strategic interaction.

This chapter highlights how these features of strategic interaction weigh heavily in real-world national security challenges in which actors' interests are neither purely aligned nor purely opposed, as is the case in undergoverned spaces, gray-zone competition, societal warfare, and global financial and public health crises.

The following sections examine the strengths and limitations of using computational complexity to quantify the complexity of strategic behavior of games and the actors within them. The first section discusses the core problem of strategic interaction in games in which actors are involved in interdependent decisionmaking processes. It then reviews the use of computational-complexity classes to analyze games, before turning to discuss factors that computational complexity obscures. Specifically, it considers such factors as multiple equilibria, communication, learning dynamics, having more than two players, and nonrational behavior. The next section shows that computational complexity is a useful notion of complexity and that there has been productive work to make the analysis more sophisticated. But it also shows that there are additional features of strategic environments that alter the complexity of analyzing and participating in such games, and these features affect the difficulty of extracting insight from a model of strategic interaction. Finally, the chapter offers some concluding thoughts and an appendix that gives a short, nontechnical introduction of game theory to support this premise.

The Core Problem of Strategic Interaction

Strategic interaction is a central tenet of the social systems and relationships between agents, whether they are individuals or groups. On a micro scale, couples with different preferences who enjoy each other's company negotiate meals, entertainment, and even parenting styles. On a macro scale, states develop trade agreements, compete in arms races, form political alliances, and conduct espionage. In all cases, the decisions and actions of one party affect the other parties; thus, all parties are anticipating and reacting to one another. This is the hallmark of a strategic interaction and interdependence.

Because strategic interaction is so pervasive, social scientists have developed tools and methods for understanding and predicting the outcomes of strategically interacting decisionmakers. Although not the only method, noncooperative game theory has emerged as a leading tool for analyzing such scenarios and anticipating potential outcomes based on the structure of the environment.[1] Stripping down problems of interdependent action in this way allows analysts to make precise statements about the structure of interactions between actors, such as characteristics of solutions, that describe and anticipate real-world events in similar situations. In its most general formulation, game theory is a formalism for analyzing decisionmakers who interact dynamically in an uncertain environment with imperfect information in pursuit of their (possibly uncertain) objectives. By combining mathematical concepts from probability theory, optimization, stochastic processes and several other disciplines, game theory provides tools to model and analyze a variety of strategic scenarios.

Definitions and Concepts

Before exploring why computational-complexity classes fail to capture the most-salient contributors to the complexity of strategic interaction as examined by game theorists, it is useful to have a core set of definitions and concepts. These definitions and concepts characterize features of strategic scenarios, players within them, and the types of solutions that can be found.

As indicated in the related chapter on gaming, the definition of a *game* is contested.[2] Game theorists also have a definition, which is, fortunately, more precise, although not necessarily

[1] Steven Tadelis, *Game Theory: An Introduction*, Princeton, N.J.: Princeton University Press, 2013.

[2] See Chapter Nineteen of this report (Elizabeth M. Bartels, Aaron B. Frank, Jasmin Léveillé, Timothy Marler, and Yuna Huh Wong, "Gaming Undergoverned Spaces: Emerging Approaches for Complex National Security Policy Problems," in Aaron B. Frank and Elizabeth M. Bartels, eds., *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches,* Santa Monica, Calif.: RAND Corporation, RR-A1275-1, 2022).

aligned with the colloquial use of the term.[3] For game theorists, a *game* is formally defined as a situation in which multiple participants (i.e., players) interact and affect each other's outcomes within a set structure.[4] The presence of two or more active players, each affecting the others, differentiates game theory from decision theory, in which a single decisionmaker chooses among options whose outcomes are unaffected by other actors.[5] Although games are commonly referred to as *strategic*, the terms *interdependent* and *contingent* are also commonly used.[6]

The outcomes of players' actions are referred to as *payoffs*, which can be distributed among the players in a variety of ways. In strictly competitive or *zero-sum* games, the payoffs are opposed and often symmetric if there are only two players, while asymmetries may exist when there are more.[7] Games in which players' payoffs reward cooperation (i.e., one's gains are not automatically the other's losses) are referred to as *general-sum* or *non-zero-sum* games and often require different approaches to solve because analysts can no longer assume that maximizing their own payoffs minimizes the payoffs of their opponents.

The concept of *equilibrium* for a game usually refers to its *Nash equilibrium*. A Nash equilibrium exists for a set of strategies in which each player cannot unilaterally improve their payoff by changing their strategy.[8] Games may have multiple equilibria, meaning that, absent any ability to coordinate their actions, players may get trapped in one equilibrium that may be less desirable to one or multiple players given alternative equilibria.[9]

Complexity in Strategic Interactions

Given a particular specification of a game, one relevant question is, "How complex is the game?" With only colloquial notions of complexity, there are several reasons why such a question is relevant. First, the complexity of a game might indicate how many resources and

[3] Garry D. Brewer and Martin Shubik, *The War Game: A Critique of Military Problem Solving*, Cambridge, Mass.: Harvard University Press, 1979, pp. 7–10.

[4] James D. Morrow, *Game Theory for Political Scientists*, Princeton, N.J.: Princeton University Press, 1994; and Éva Tardos and Vijay V. Vazirani, "Basic Solution Concepts and Computational Issues," in Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, eds., *Algorithmic Game Theory*, New York: Cambridge University Press, 2007, p. 3.

[5] George Tsebelis, "The Abuse of Probability in Political Analysis: The Robinson Crusoe Fallacy," *American Political Science Review*, Vol. 83, No. 1, March 1989.

[6] Arthur A. Stein, *Why Nations Cooperate: Circumstance and Choice in International Relations*, Ithaca, N.Y.: Cornell University Press, 1990.

[7] Martin J. Osborne, *An Introduction to Game Theory*, 1st ed., New York: Oxford University Press, 2003, pp. 365–366.

[8] Osborne, 2003, pp. 21–52.

[9] Nolan McCarty and Adam Meirowitz, *Political Game Theory: An Introduction*, New York: Cambridge University Press, 2007, p. 254.

how much effort decisionmakers might have to expend to participate in it. The more complex the game, the more resources participants must expend to determine their optimal strategies. This can foreshadow barriers to entry into certain domains for which individuals do not have the resources to fully understand or compete in the strategic environment. Such problems may include decisions by firms to enter into new markets or by states to enter into an arms race or alliance-building to balance against rivals.[10]

Second, and relatedly, as the complexity of a game increases, the likelihood that decisionmakers will make optimal choices decreases, given what the other players are doing.[11] Instead, they will use rules, heuristics, and *satisficing behavior*—often affected (whether positively or negatively) by cognitive biases—to cope with the complexity, and the ultimate outcome might differ from what a set of optimizing decisionmakers would achieve. This dynamic means that a complex game environment adds to the complexity of individual agents, which, in turn, makes the entire game environment more complex. Ultimately, understanding the complexity of a game can provide clues as to how agents might differ from purely rational and optimizing agents.

Third, as strategic decisionmaking shifts from human decisionmakers to automated algorithmic decisionmakers, complexity will partly determine how well those algorithms perform, the predictability of their decisions, and their vulnerability to adversarial attacks.[12] It will also inform how much a machine needs to be trained, tested, and evaluated before being deployed in an environment. Therefore, understanding the complexity of a strategic interaction is a necessary condition for the successful automation and improvement of human decisionmaking by machines.

Given the relative importance of complexity, it is no surprise that quantitative measures of strategic complexity have emerged. Nowhere has this been more apparent than in the notion of computational-complexity theory.[13] However, complexity classes are meant to apply to a variety of problems, and features of any specific problem are abstracted away in pursuit of generality. For this reason, the usual classifications of computational complexity do not give a complete picture of the complexity of games and the strategies of the actors in them. Such characteristics as multiple equilibria, communication, chaotic learning dynamics, alternative solution concepts, and nonrational behavior all contribute to the complexity of understand-

[10] Robert Jervis, "Cooperation Under the Security Dilemma," *World Politics*, Vol. 30, No. 2, January 1978; Andrew Kydd, "Game Theory and the Spiral Model," *World Politics*, Vol. 49, No. 3, April 1997; F. Warren McFarlan, "Information Technology Changes the Way You Compete," *Harvard Business Review*, Vol. 62, May 1984; and Alastair Smith, "Alliance Formation and War," *International Studies Quarterly*, Vol. 39, No. 4, December 1995.

[11] Sendhil Mullainathan and Richard H. Thaler, "Behavioral Economics," Cambridge, Mass.: National Bureau of Economic Research, NBER Working Paper 7948, October 2000.

[12] In this context, *adversarial attacks* are attacks that manipulate inputs to algorithms to affect outputs of computations.

[13] Sanjeev Arora and Boaz Barak, *Computational Complexity: A Modern Approach*, New York: Cambridge University Press, 2009.

ing and participating in a strategic environment. Furthermore, the importance of these characteristics is amplified in games in which players' goals are neither completely opposed nor completely aligned (i.e., general-sum games).

Computational-Complexity Classes to Analyze Games

At its core, the theory of computational complexity tries to quantify how difficult it is to solve a problem with a given set of characteristics. Although the rigorous definition of *difficult* requires a full treatment of Turing machines (an abstract mathematical model of a computer), an informal treatment is sufficient to understand computational complexity and its application to games.

One of the key questions in computational-complexity theory is whether questions whose potential answers can be easily verified can also be easily solved. As a canonical example, consider the set of numbers

$$S = \{-12, -8, -1, 2, 5, 4\}$$

and the subset-sum question, "Does there exist a subset of S that sums to 0?" Of course, if someone proposes an answer—say, $A_1 = \{2,5,-8\}$—it is easy to verify that 2 + 5 − 8 does not equal 0. However, it is not as straightforward to find a way to *determine* whether there is a subset of numbers that sums to 0. One algorithm for doing so is simply to try all possible subsets, which, for the set S consisting of eight elements, would involve 64 such subsets. For a set of 20 numbers, a search of all possible subsets would involve more than 1,000,000 different permutations. Thus, an algorithm seeking to search across all subsets of a set of numbers scales poorly with the size of the set.[14] This notion of verifiability versus solvability is the central concept behind the classic "P vs. NP" problem at the core of computational-complexity class analysis.[15]

Computational Complexity and Games

At this point, it is unclear how these abstract concepts of computational complexity apply to games. To draw the connection, consider the following Nash existence question: "Given a game with a set of players, strategies, and payoffs, does there exist a Nash equilibrium?" This question is parallel to the question in the previous section. The general format is, "Given an object, does that object have a certain property?" In the subset-sum problem, the object is the set S, and the property is "has a subset that sums to 0." In the Nash existence question, the

[14] 5 + 4 − 8 − 1 = 0.

[15] Michael Sipser, "The History and Status of the P Versus NP Question," *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing*, New York: Association for Computing Machinery, July 1992.

object is a game, and the property is "has a Nash equilibrium." Given the parallel nature of these questions, one might be tempted to ask, "How hard is it to answer the Nash existence question as the number of players and strategies grows?"

At a first pass, the answer to whether a game possesses a Nash equilibrium is trivial because Nash's famous theorem proved that *every* finite game has at least one Nash equilibrium.[16] However, a slightly different question is also of interest: "Given a game, what *is* the Nash equilibrium?" This "find-the-Nash" question has demanded much attention because it is a special type of problem for which a solution is known to exist but for which it is not clear how difficult the solution is to find. It turns out that the find-the-Nash problem has similar properties as the subset-sum problem. Specifically, it is relatively easy to verify whether a particular set of strategies forms a Nash equilibrium, but the difficulty of *finding* an equilibrium strategy (with any known algorithm) increases exponentially with the number of strategies.[17] For this reason, finding a Nash equilibrium is known to be computationally difficult.[18]

This foundational result establishes that finding a Nash equilibrium is inherently a difficult problem and has been extended to prove other game-theoretic insights. For example, the following questions are also computationally hard:[19]

- Does a second Nash equilibrium exist?
- What is the socially optimal Nash equilibrium?
- Does there exist a Nash equilibrium where one player's payoff is at least x?

These results also have implications for the complexity of dynamic games and games of imperfect information. Broadly speaking, introducing dynamics or imperfect information to a static game of perfect information tends to increase the number of strategies for each player. This is because a strategy specifies what a player would do at *each* point in time and for any amount of imperfect information. The computational-complexity theory has established that the difficulty of computing an equilibrium increases exponentially given the number of strategies, so adding dynamic components with imperfect information to a static game does increase the game's computational complexity.

One overarching feature of the computational-complexity insights is that they make minimal assumptions about the underlying game. This generality can obfuscate the complexity of finding a Nash equilibrium for a particular game with known structure. For example, if

[16] John Nash, "Non-Cooperative Games," *Annals of Mathematics*, Vol. 54, No. 2, September 1951.

[17] Christos H. Papadimitriou, "The Complexity of Finding Nash Equilibria," in Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, eds., *Algorithmic Game Theory*, New York: Cambridge University Press, 2007, p. 30.

[18] Technically, "find-the-Nash" is in PPAD complete, a complex class that stands for *Polynomial Parity Argument* (Directed Case) (Papadimitriou, 2007, p. 38).

[19] Vincent Conitzer and Tuomas Sandholm, *Complexity Results About Nash Equilibria*, arXiv.org, May 28, 2002.

the game is a two-player zero-sum game, it is not computationally difficult to find a Nash equilibrium, because the game can be solved with a linear program.[20] In this case, the characterization in terms of computational complexity *overstates* the difficulty of finding an equilibrium. Then again, even though it is not computationally hard to find an equilibrium, the problem still might be intractable given the game's underlying structure. For example, the game of Go is a two-player zero-sum game, a class of game for which finding a Nash equilibrium is not computationally hard. However, because the number of strategies is, famously, greater than the estimated number of atoms in the universe, it is believed to be impossible to find an equilibrium of the game—a provable theorem that demonstrates that, if both players make optimal decisions during gameplay, the outcome of the game is knowable from the first move, with the first or second mover winning or the game ending in a draw.[21] Although the artificial intelligence program AlphaGo claimed victories over Lee Sodol, the highest-ranked Go player in the world, it subsequently lost to future generations of its own program, AlphaZero, demonstrating that potentially novel and even-higher-performing strategies remain to be discovered and that a gap exists between superhuman performance and optimal play.[22] Thus, in the case of a relatively simple game with a large state space, computational-complexity results *understate* the difficulty of finding an equilibrium.

All told, complexity class analysis provides a general framework for analyzing the difficulty of computing a Nash equilibrium in a game. Of course, if players try to play equilibrium strategies, then such analysis quantifies how difficult it is for players to find an optimal strategy. However, in the pursuit of generality, computational-complexity analysis may obfuscate the true difficulty in participating in or analyzing a strategic interaction. As a result, understanding the underlying structure of a particular game is crucial for refining the notion of complexity in strategic interaction.

For a specific example, consider a competitive resource allocation problem in which two military organizations must allocate disparate resources (e.g., intelligence, combat, logistics, and command and control assets) across multiple distinct battlefields, each organization trying to maximize its chances of defeating its rival on as many battlefields as possible. This

[20] Conitzer and Sandholm, 2002.

[21] John Tromp, "The Number of Legal Go Positions," in Aske Plaat, Walter Kosters, and Jaap van den Herik, eds., *Computers and Games*, Cham, Switzerland: Springer International Publishing, 2016.

[22] David Silver and Demis Hassabis, "AlphaGo: Mastering the Ancient Game of Go with Machine Learning," *Google AI Blog*, January 27, 2016; David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al., "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature*, Vol. 529, No. 7587, January 28, 2016; David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al., *Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm*, ArXiv.org, December 5, 2017; and David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al., "A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go Through Self-Play," *Science*, Vol. 362, No. 6419, December 7, 2018.

game is known as the *multiresource Blotto game*, and, given its zero-sum nature, finding a solution to this game is technically less complex than finding an equilibrium in a game in which the players' goals are not perfectly opposed.[23] However, because the number of strategies grows exponentially with the number of different resources, solving the game with more than five or six different resource types quickly becomes intractable. This illustrates how even games that are not technically complex from a structural perspective can nevertheless become difficult to solve.

Factors That Computational-Complexity Theory Obscures

One of the most prominent characteristics of a game is whether it is competitive or cooperative in nature. The most competitive games are known as *zero sum* and encompass many parlor games, resource allocation games, and attacker-defender games.[24] Cooperative games are the opposite; all players share the same payoff; thus, what is good for one player is good for the other players. Common examples of cooperative games are coordinating aircraft to fight a wildfire, designing optimal organizational communication protocols, and organizing social movements.[25] Of course, there are games that are neither purely competitive nor purely cooperative, such as auctions, political negotiations and lobbying, firm price setting, deterrence, persuasion, and personnel decisions.[26] These cases of mixed cooperation and competition add complexity to understanding and participating in a game.

---

[23] Soheil Behnezhad, Sina Dehghani, Mahsa Derakhshan, Mohammad Taghi Haji Aghayi, and Saeed Seddighin, "Faster and Simpler Algorithm for Optimal Strategies of Blotto Game," *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, San Francisco, Calif.: AAAI Press, 2017.

[24] D. W. Blackett, "Some Blotto Games," *Naval Research Logistics Quarterly*, Vol. 1, No. 1, 1954; and E. O. Ibidunmoye, B. K. Alese, and O. S. Ogundele, "Modeling Attacker-Defender Interaction as a Zero-Sum Stochastic Game," *Journal of Computer Sciences and Applications*, Vol. 1, No. 2, 2013.

[25] John Ginkel and Alastair Smith, "So You Say You Want a Revolution: A Game Theoretic Explanation of Revolution in Repressive Regimes," *Journal of Conflict Resolution*, Vol. 43, No. 3, June 1999; Peter W. Kennedy, "Information Processing and Organization Design," *Journal of Economic Behavior & Organization*, Vol. 25, No. 1, 1994; and Kunal Menda, Yi-Chun Chen, Justin Grana, James W. Bono, Brendan D. Tracey, Mykel J. Kochenderfer, and David Wolpert, "Deep Reinforcement Learning for Event-Driven Multi-Agent Decision Processes," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 20, No. 4, April 2019.

[26] Fred Charles Iklé and Nathan Leites, "Political Negotiation as a Process of Modifying Utilities," *Journal of Conflict Resolution*, Vol. 6, No. 1, 1962; Emir Kamenica and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, Vol. 101, No. 6, October 2011; Vijay Krishna, *Auction Theory*, 2nd ed., Burlington, Mass.: Elsevier, Academic Press, 2009; Andrew Ledvina and Ronnie Sircar, "Dynamic Bertrand Oligopoly," *Applied Mathematics & Optimization*, Vol. 63, No. 1, 2011; Michael Spence, "Job Market Signaling," *Quarterly Journal of Economics,* Vol. 87, No. 3, August 1973; and Franz Wirl, "The Dynamics of Lobbying—A Differential Game," *Public Choice*, Vol. 80, No. 3, 1994.

Strategic Communication

Communication—sometimes referred to as *signaling*—plays a key role in many strategic scenarios. A potential employee communicating their value to a company, a lawyer choosing evidence to reveal to a jury, and a world power promising retaliation for cyberattacks provide examples in which communication is central to strategic interaction. However, communication is not always verifiable, is sometimes noisy, and can sometimes be misinterpreted. For these reasons, including communication as part of an agent's strategy can add to the complexity of a game.

However, in two-player zero-sum games, it is always optimal to ignore any communication from the other player.[27] The reason is that, if the players' interests are diametrically opposed, neither player has an incentive to convey valuable information to the other. Therefore, any communication must be meaningless and can be ignored. This means that two-player zero-sum games are not riddled with the complexity that arises from strategic communication. However, in common-interest games, truthful communication can only be helpful. Specifically, if there is infinite communication ability, the optimal communication protocol for a player is simple: Truthfully reveal all information.[28] The intuition is that, if all players have the exact same goal, no player has an incentive to deceive any other player and, thus, all communication is truthful and the receiver of any piece of communication can be sure that the information is true.

Importantly, in games in which interests are neither completely aligned nor completely opposed, communication can become complex. In such scenarios, some players might have an incentive to reveal partial information. Furthermore, there might be an incentive for players to tell the truth only if certain other conditions are met. For example, two firms under price competition might be willing to truthfully share market research information *only if* they are sufficiently forward-looking. Otherwise, there is an incentive for one firm to deceive the other, and communication breaks down.[29] Other examples of complex communication in games are job market signaling, product quality signaling, and deterrence.[30] The analytical problems posed by communication between players are exacerbated if the differential costs and modes of signaling are considered, such as those costs that are paid immediately, through actions (e.g., the movement of military forces into a theater), versus those that might be paid

---

[27] Joseph Farrell, "Talk Is Cheap," *American Economic Review*, Vol. 85, No. 2, May 1995.

[28] Jasmina Arifovic and B. Curtis Eaton, "The Evolution of Type Communication in a Sender/Receiver Game of Common Interest with Cheap Talk," *Journal of Economic Dynamics and Control*, Vol. 22, Nos. 8–9, 1998.

[29] Justin Grana, James Bono, and David Wolpert, "Reasoning About 'When' Instead of 'What': Collusive Equilibria with Stochastic Timing in Repeated Oligopoly," *B.E. Journal of Theoretical Economics*, Vol. 20, No. 1, January 2020.

[30] Maarten C. W. Janssen and Santanu Roy, "Signaling Quality Through Prices in an Oligopoly," *Games and Economic Behavior*, Vol. 68, No. 1, January 2010; Spence, 1973; and Jonathan William Welburn, Justin Grana, and Karen Schwindt, "Cyber Deterrence or: How We Learned to Stop Worrying and Love the Signal," Santa Monica, Calif.: RAND Corporation, WR-1294-OSD, 2019.

only if the recipient does not do as desired (e.g., the declaration of redlines regarding the use of military force).[31]

In short, allowing for communication in a game can increase its complexity and richness of outcomes. However, this richness is most pronounced in games in which agents' incentives are only partially aligned.

From a security perspective, consider interactions in an undergoverned space. Suppose one party is trying to restore order but, to do so, must gather information from several informants, each with a unique preference on how order should be restored. From the perspective of the informants, each does not have the incentive to fully tell the truth, because each is trying to influence the future order. Furthermore, each informant will form beliefs about what other informants are saying and might try to strategically align or oppose their message with the messages of other informants. The party trying to restore order must weigh each piece of information knowing that each informant is trying to influence its actions. It is often not optimal for the order-restoring party to ignore *all* information, because some of it is true, but it must carefully reason about how the informants might be incentivized to lie and ultimately make decisions based on unreliable information. A similar logic plays out when mediators seek to prevent the occurrence of a conflict or its further escalation and conflicting parties seek to assess not only the resolve of one another but also the bias and credibility of the mediator.[32]

Equilibrium Selection

Another feature of strategic interaction that can add to the complexity of predicting and understanding human behavior is the existence of multiple equilibria.[33] In practical terms, there might be several outcomes of a strategic interaction that are consistent with common notions of boundedly rational behavior. However, it is not necessarily clear which of those outcomes will be realized.

The simplest example of multiple equilibria is in coordination games. Consider two players who made plans to meet at a restaurant but forgot where to meet and cannot communicate. However, they know that they were going to meet at either the town's pizza parlor or its steakhouse. The players are better off if they choose the same restaurant. In this case, there are multiple equilibria: both players choosing to go to the pizza parlor, and both players choosing to go to the steakhouse.[34] Although both outcomes are plausible, it would not be clear to an external observer how to choose between the two outcomes to predict where the

[31] James D. Fearon, "Signaling Foreign Policy Interests: Tying Hands Versus Sinking Costs," *Journal of Conflict Resolution*, Vol. 41, No. 1, February 1997.

[32] Andrew Kydd, "Which Side Are You On? Bias, Credibility, and Mediation," *American Journal of Political Science*, Vol. 47, No. 4, October 2003.

[33] John C. Harsanyi and Reinhard Selten, *A General Theory of Equilibrium Selection in Games*, Cambridge, Mass.: MIT Press, 1988.

[34] There is also a mixed-strategy equilibrium.

diners will meet. Examples from a security perspective are coordinated hardening of critical infrastructure, coordinated responses among allies to terrorist aggression, and coordinated responses to the release of a biological agent.[35] In all of these cases, there might be many possible solutions, but it is most important that several parties coordinate on a plan; which plan they choose is less important.

Again, the complexities of multiple equilibria are most pronounced in games in which players' interests are neither perfectly aligned nor perfectly opposed. For example, in two-player zero-sum games, all equilibria give the same payoff to the players; the equilibria are *payoff equivalent*. Therefore, choosing among equilibria becomes less important because the welfare of each player is constant among equilibria. This is especially poignant in zero-sum games because, if a player wins in one equilibrium, that player will win in all equilibria. Then again, when players' interests are perfectly aligned, there always exists an equilibrium in which no collective action of any number of players can make the team better off.[36]

Unfortunately, neither of these properties holds for general-sum games. That is, all equilibria do not necessarily yield the same payoffs, and there does not always exist a Pareto-Optimal equilibrium. As a result, it is unclear how to predict the outcome of a strategic interaction in general-sum games. As a concrete illustration, consider a slight tweak to the restaurant coordination game. The players still want to go to the same restaurant, but now assume that player 1 slightly prefers pizza over steak while player 2 slightly prefers steak over pizza. Would player 1 go to the pizza parlor because they prefer pizza, or would player 1 go to the steakhouse because they know player 2 prefers steak? How would player 2 act with a similar line of reasoning? Without further structure on the strategic interaction, it is not clear how to choose among the possible reasonable outcomes.

Learning Dynamics

Although many questions about strategic interaction are concerned with predicting ultimate outcomes, it is sometimes relevant to understand how players arrive at certain outcomes and how they adjust their decisions based on past observations.[37] How players learn and possibly converge to a stable outcome is often as important as the outcome itself. A common argument is that, if a relatively simple learning procedure does converge to an equilibrium outcome, then there is evidence that human learners would converge to the same outcome.[38]

---

35 Todd Sandler, "Collective Versus Unilateral Responses to Terrorism," *Public Choice*, Vol. 124, No. 1, July 2005.

36 Katja Verbeeck, Ann Nowé, Tom Lenaerts, and Johan Parent, "Learning to Reach the Pareto Optimal Nash Equilibrium as a Team," in Bob McKay and John Slaney, eds., *AI 2002: Advances in Artificial Intelligence*, Berlin, Germany: Springer, 2002.

37 Robert Powell, "Bargaining and Learning While Fighting," *American Journal of Political Science*, Vol. 48, No. 2, April 2004.

38 Drew Fudenberg and David K. Levine, *The Theory of Learning in Games*, reprint ed., Cambridge, Mass.: MIT Press, 1998.

Once again, the complexities that arise when analyzing learning dynamics are most pronounced in general-sum games. In two-player zero-sum games, a simple algorithm known as *fictious play* yields a result equivalent to the Nash equilibrium strategies.[39] This algorithm also applies to games of common interest. Furthermore, an algorithm known as *iterated best response*, in which the player simply best responds to their opponent's previous action, is known to converge in team games. However, these theoretical convergence guarantees do *not* extend to general-sum games. It is easy to see that fictious play leads to chaotic learning dynamics in a slightly modified game of rock paper scissors.[40] Nuances in learning dynamics suggest that analyzing how players might learn and adapt their strategies is much more complex in general-sum games than in two-player zero-sum games or team games.

More Than Two Players

As the previous subsection noted, common issues that make games complex occur when players are neither purely competitive nor purely cooperative. However, many of the reductions in complexity that occur when two players have perfectly opposed objectives cease to be viable when a game contains more than two active players. For example, suppose there are three opposing states that want to conquer the territory of the others, but the cost of going to war with any rival is not worth the value of the territorial gains that victory would bring. However, if two states declare war against the third, the solo country immediately forfeits its land because it knows that it cannot win against the combined forces of the other two, allowing the victors to divide the surrendered territory evenly without paying the costs of war. In this game, the equilibrium is for any two countries to ally against one country. However, standard game-theoretic arguments do not identify *which* of the two countries should form an alliance; thus, the problem of multiple equilibria exists despite the game being zero sum.[41]

Bounded Rationality

Although the previous results showed how complexities from communication, multiple equilibria, and learning dynamics are most pronounced in general-sum games, the complexities from bounded rationality are part of all types of games. In brief, psychological experiments have shown that human decisionmakers are subject to cognitive computational constraints, such as limited attention and working memory.[42] As a result, decisionmakers deviate from

[39] Osborne, 2003, pp. 136–137.

[40] Yuzuru Sato, Eizo Akiyama, and J. Doyne Farmer, "Chaos in Learning a Simple Two-Person Game," *Proceedings of the National Academy of Sciences*, Vol. 99, No. 7, April 2, 2002.

[41] Bruce M. Russett, "Components of an Operational Theory of International Alliance Formation," *Journal of Conflict Resolution*, Vol. 12, No. 3, September 1, 1968.

[42] Daniel Kahneman, *Thinking, Fast and Slow*, New York: Farrar, Straus and Giroux, 2013; George A. Miller, "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing

the expectations of full rationality and optimal utility maximization, instead engaging in satisficing behavior using a variety of heuristics.[43]

Although deviations from rationality are ubiquitous, there are few guidelines for predicting in which strategic situations real human decisionmakers will act with predictable systematic biases. For example, consider two different types of documented biases in decisionmaking: anchoring and base rate neglect. In anchoring, an individual's initial beliefs about an uncertain entity or likelihood of an event are not sufficiently updated, despite the availability of new information.[44] By contrast, base rate neglect biases occur when individuals overreact to new information as it is presented, undervaluing the larger sample provided by history.[45] These biases work in opposite directions; thus, someone trying to *predict* the outcome of a strategic interaction would benefit from knowing which of these two (or other) behavioral biases might be exhibited by players. Unfortunately, it is difficult and perhaps even impossible to predict which biases people will exhibit before the biases are observed.

Further compounding issues, human decisionmakers are more likely to exhibit behavioral biases when the strategic environment is already complex.[46] Stochastic dynamics, high-dimensional strategy and observation spaces, and complex communication protocols can all instigate deviations from full rationality. Because nonrational behavior is complex in itself, this is a case in which the effects of complexity are compounded. General-sum games with communication and multiple equilibria are complex, even when it is assumed that players are capable of making fully rational, optimal decisions. However, a complex strategic environment encourages behavioral nonrational decisionmaking, which increases the complexity of analyzing or predicting the behavior of decisionmakers.

Information," *Psychological Review*, Vol. 63, No. 2, March 1956; and Herbert A. Simon, "Bounded Rationality in Social Science: Today and Tomorrow," *Mind & Society*, Vol. 1, No. 1, March 2000.

43 Robert A. Becker and Subir K. Chakrabarti, "Satisficing Behavior, Brouwer's Fixed Point Theorem and Nash Equilibrium," *Economic Theory*, Vol. 26, No. 1, July 2005; Hunter Crowther-Heyck,. *Herbert A. Simon: The Bounds of Reason in Modern America*, Baltimore, Md.: Johns Hopkins University Press, 2005; Gerd Gigerenzer, Ralph Hertwig, and Thorsten Pachur, eds., *Heuristics: The Foundations of Adaptive Behavior*, New York: Oxford University Press, 2015; Thomas Gilovich, Dale Griffin, and Daniel Kahneman, eds., *Heuristics and Biases: The Psychology of Intuitive Judgment*, New York: Cambridge University Press, 2002; and Gary Klein, *Sources of Power: How People Make Decisions*, Cambridge, Mass.: MIT Press, 1999.

44 Jørgen Vitting Andersen, "Detecting Anchoring in Financial Markets," *Journal of Behavioral Finance*, Vol. 11, No. 2, 2010; and Robert Jervis, *How Statesmen Think: The Psychology of International Politics,* Princeton, N.J.: Princeton University Press, 2017.

45 Tilmann Betsch, Glenn-Merten Biel, Claudia Eddelbüttel, and Andreas Mock, "Natural Sampling and Base-Rate Neglect," *European Journal of Social Psychology*, Vol. 28, No. 2, March–April 1998.

46 Mullainathan and Thaler, 2000.

Benefit of Computational-Complexity Theory: Quantifying Complexity

One benefit of computational-complexity theory is that it provides a specific quantitative measure about the computational resources needed to compute a solution to a game. This number—usually a function of the strategies—is concrete and based on mathematical principles. Unfortunately, some of the sources of complexity discussed earlier are not amenable to such a rich classification, and such quantification based on computational complexity may obfuscate the sources of strategic complexity that characterize the specific game and the strategies of the players being analyzed. For example, computational complexity might say how long it takes to find an equilibrium, but it does not address how to choose among multiple equilibria. However, there are binary measures that can be used to quantify complexity.

The first, and arguably most important, binary measure is whether agents' goals are either perfectly aligned or perfectly opposed. As discussed in the previous section, there are several nuanced considerations that arise when agents have some incentive to collaborate but also to compete.

The second binary measure is whether players communicate. In games in which players are perfectly aligned, honest communication is optimal, but designing communication protocols can be difficult. When players' goals are not perfectly aligned, players might be incentivized to tell partial truths, balancing the benefits of cooperation with the advantages derived from information asymmetries. Both players deciding how to communicate a partial truth and how to disentangle facts from falsehoods adds complexity to strategic interaction.

The third binary measure is whether players are experts interacting in a familiar domain or neophytes interacting in a novel domain. In the latter case, it is unlikely that players will act rationally, and it is likely that they will deviate from optimal behavior in systematically identifiable yet unpredictable ways. Crucially, these deviations do not just add noise to potential outcomes but can drastically alter the distribution of outcomes, increasing the complexity of analyzing and participating in strategic games.

Unfortunately, there is no principled way to combine these binary measures into one number that captures the complexity of the game. For this reason, a careful analysis of each strategic interaction is required to understand the effects of each feature that increases complexity.

Concluding Thoughts

The computational complexity of finding a Nash equilibrium provides valuable insight into one dimension of complexity in a game. The results in computational-complexity theory are general and apply to broad classes of games. However, there are other features of strategic interactions that alter a game's complexity on other dimensions. Specifically, such concepts as communication, multiple equilibria, and learning dynamics can all add to the complexity of understanding, predicting, and participating in a strategic interaction. These features are most pronounced in general-sum games in which players' interests are neither directly opposed nor aligned. This suggests that, although recent advances in computing strategies in two-player zero-sum games, such as Go, can handle the computational complexity that arises from a large strategy space, there are other sources of complexity in real-world general-sum strategic interaction that require additional treatment and analysis.

Bounded rationality adds another layer of complexity. Although the rational-actor model assumes optimal behavior in pursuit of internally consistent utilities, there are many reasons that players can deviate from this expectation. They can have error-prone strategies, engage in limited hierarchical thinking, or systematically miscalculate probability and rewards. However, when one is reasoning and making predictions about real humans in real strategic situations, it is often difficult to predict which, if any, of the systematic errors they would make. Because behavioral and systematic errors are well documented in real-world decisionmakers, this is another source of complexity in understanding human decisionmaking that computational-complexity theory overlooks.

Appendix: Game Theory—From Normal Form to Imperfect Information in Continuous Time

Game theory is the study of multiple interacting decisionmakers (players) set in the (boundedly) rational-actor context. Specifically, decisionmakers have strategies, payoffs, *information sets* (specifications of what each player knows), and chance elements. These basic building blocks provide a rich set of tools to *specify* a game. Once a game is specified, a theorist employs a *solution concept* to draw conclusions on how the game's building blocks influence behavior. Although the Nash equilibrium solution concept is the most well known, its refinements, such as subgame perfect equilibrium, perfect Bayesian equilibrium, sequential equilibrium, and forward induction equilibrium, are all commonly used by game theorists. In addition, behavioral solution concepts, such as the quantal response equilibrium, the level-k solution concept, and prospect theory preferences, add additional flexibility into predicting human behavior.

The simplest version of a game is the canonical simultaneous-move game. In this game, players act only once in a deterministic environment and choose their action among a finite strategy set. These types of games are often used to introduce such concepts as dominant strategies, coordination problems, and mixed strategies. Although such simultaneous-move games capture some real-world decision problems (mismatch games, such as those between a penalty kicker and a goalie, are a common example), the single-shot simultaneous-move game lacks many of the key elements that makes real-world strategic problems complex.

Beyond simultaneous-move games are dynamic games of perfect information. In these games, players take actions at discrete moments in time. Importantly, one of the players can be a "nature" player that represents the underlying environment; thus, dynamic games can capture future uncertainty. Many parlor games, such as chess, checkers, Go, and backgammon, are dynamic games of perfect information. Although all of these games have only two players that take turns, turn-taking is not a restriction of dynamic games of perfect information; instead, the formalism allows for any (possibly random) specification of move order. Similarly, at some instants, the players can make moves simultaneously.

Although identical in structure to dynamic games of perfect information, Markov games provide an intuitive framework for analyzing dynamic games of perfect information when the nature player has a significant role. In Markov games, the nature player can be used to model the evolution of the parameters of the game. For example, the rules of the game can change, the players' preferences can change, the number of players can change, and the available strategies of each player can change. Furthermore, these can be either random *or* determined by players' past actions. Although the state change must be *Markovian*—meaning that transition probabilities to future states are determined only by the game's current state—a careful definition of the *state space* allows for long-term dependencies in the game's state-space transitions to be considered. The discrete-time Markovian framework can be extended to continuous time; players act in discrete intervals, and the time between moves is governed by a stochastic clock whose distribution is governed by chance and players' actions. These are known as *semi-Markov games* and are the most general formulation of dynamic games of perfect information with discrete action times.

Of course, one central feature of real-world strategic decision problems is imperfect information. In the simplest sense, a game of poker has imperfect information because each player does not know the other players' cards. However, many real-world strategic decision problems—auctions, hiring decisions, pricing decisions, investment and finance, arms races, lobbying, and almost all of international relations—are rife with imperfect information.[47] The game-theoretic formalism is rich enough to insert imperfect information into both simultaneous-move and dynamic games. In these cases, players must form beliefs about the true state of the world and take actions given their beliefs. In a general sense, semi-Markov games of imperfect information allow for environments in which the underlying rules and preferences change over time but players are not fully aware of these changes. In such scenarios, players may receive heterogeneous information asynchronously and take hidden, perfectly observable, or partially observable actions.

In total, a semi-Markov game of imperfect information (which subsumes many simpler formulations) is flexible enough to capture, in theory, many real-world strategic scenarios. How appropriate such an approach might be to understanding a given strategic scenario is less clear. To address this question of appropriateness, other conditions, such as tractability and model sophistication, are important. These features are often subjective and are left to the judgment of the modeler. So, although game theory is a flexible mathematical tool that

[47] James D. Fearon, "Rationalist Explanations for War," *International Organization*, Vol. 49, No. 3, Summer 1995; and Kristopher W. Ramsay, "Information, Uncertainty, and War," *Annual Review of Political Science*, Vol. 20, No. 1, 2017.

can capture an array of strategic interactions, the appropriateness of such an approach, like all models, is often subjective and will provide some insights while obscuring others.

Acknowledgments

I would like to thank Emily Ward, who edited the chapter.

References

Andersen, Jørgen Vitting, "Detecting Anchoring in Financial Markets," *Journal of Behavioral Finance*, Vol. 11, No. 2, 2010, pp. 129–133.

Arifovic, Jasmina, and B. Curtis Eaton, "The Evolution of Type Communication in a Sender/Receiver Game of Common Interest with Cheap Talk," *Journal of Economic Dynamics and Control*, Vol. 22, Nos. 8–9, 1998, pp. 1187–1207.

Arora, Sanjeev, and Boaz Barak, *Computational Complexity: A Modern Approach*, New York: Cambridge University Press, 2009.

Bartels, Elizabeth M., Aaron B. Frank, Jasmin Léveillé, Timothy Marler, and Yuna Huh Wong, "Gaming Undergoverned Spaces: Emerging Approaches for Complex National Security Policy Problems," in Aaron B. Frank and Elizabeth M. Bartels, eds., *Adaptive Engagement for Undergoverned Spaces: Concepts, Challenges, and Prospects for New Approaches,* Santa Monica, Calif.: RAND Corporation, RR-A1275-1, 2022. As of July 2022: https://www.rand.org/pubs/research_reports/RRA1275-1.html

Becker, Robert A., and Subir K. Chakrabarti, "Satisficing Behavior, Brouwer's Fixed Point Theorem and Nash Equilibrium," *Economic Theory*, Vol. 26, No. 1, July 2005, pp. 63–83.

Behnezhad, Soheil, Sina Dehghani, Mahsa Derakhshan, Mohammad Taghi Haji Aghayi, and Saeed Seddighin, "Faster and Simpler Algorithm for Optimal Strategies of Blotto Game," *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, San Francisco, Calif.: AAAI Press, 2017, pp. 369–375.

Betsch, Tilmann, Glenn-Merten Biel, Claudia Eddelbüttel, and Andreas Mock, "Natural Sampling and Base-Rate Neglect," *European Journal of Social Psychology*, Vol. 28, No. 2, March–April 1998, pp. 269–273.

Blackett, D. W., "Some Blotto Games," *Naval Research Logistics Quarterly*, Vol. 1, No. 1, 1954, pp. 55–60.

Brewer, Garry D., and Martin Shubik, *The War Game: A Critique of Military Problem Solving*, Cambridge, Mass.: Harvard University Press, 1979.

Conitzer, Vincent, and Tuomas Sandholm, *Complexity Results About Nash Equilibria*, arXiv.org, May 28, 2002.

Crowther-Heyck, Hunter, *Herbert A. Simon: The Bounds of Reason in Modern America*, Baltimore, Md.: Johns Hopkins University Press, 2005.

Farrell, Joseph, "Talk Is Cheap," *American Economic Review*, Vol. 85, No. 2, May 1995, pp. 186–190.

Fearon, James D., "Rationalist Explanations for War," *International Organization*, Vol. 49, No. 3, Summer 1995, pp. 379–414.

Fearon, James D., "Signaling Foreign Policy Interests: Tying Hands Versus Sinking Costs," *Journal of Conflict Resolution*, Vol. 41, No. 1, February 1997, pp. 68–90.

Fudenberg, Drew, and David K. Levine, *The Theory of Learning in Games*, reprint ed., Cambridge, Mass.: MIT Press, 1998.

Gigerenzer, Gerd, Ralph Hertwig, and Thorsten Pachur, eds., *Heuristics: The Foundations of Adaptive Behavior*, New York: Oxford University Press, 2015.

Gilovich, Thomas, Dale Griffin, and Daniel Kahneman, eds., *Heuristics and Biases: The Psychology of Intuitive Judgment*, New York: Cambridge University Press, 2002.

Ginkel, John, and Alastair Smith, "So You Say You Want a Revolution: A Game Theoretic Explanation of Revolution in Repressive Regimes," *Journal of Conflict Resolution*, Vol. 43, No. 3, June 1999, pp. 291–316.

Grana, Justin, James Bono, and David Wolpert, "Reasoning About 'When' Instead of 'What': Collusive Equilibria with Stochastic Timing in Repeated Oligopoly," *B.E. Journal of Theoretical Economics*, Vol. 20, No. 1, January 2020.

Harsanyi, John C., and Reinhard Selten, *A General Theory of Equilibrium Selection in Games*, Cambridge, Mass.: MIT Press, 1988.

Ibidunmoye, E. O., B. K. Alese, and O. S. Ogundele, "Modeling Attacker-Defender Interaction as a Zero-Sum Stochastic Game," *Journal of Computer Sciences and Applications*, Vol. 1, No. 2, 2013, pp. 27–32.

Iklé, Fred Charles, and Nathan Leites, "Political Negotiation as a Process of Modifying Utilities," *Journal of Conflict Resolution*, Vol. 6, No. 1, 1962, pp. 19–28.

Janssen, Maarten C. W., and Santanu Roy, "Signaling Quality Through Prices in an Oligopoly," *Games and Economic Behavior*, Vol. 68, No. 1, January 2010, pp. 192–207.

Jervis, Robert, "Cooperation Under the Security Dilemma," *World Politics*, Vol. 30, No. 2, January 1978, pp. 167–214.

Jervis, Robert, *How Statesmen Think: The Psychology of International Politics*, Princeton, N.J.: Princeton University Press, 2017.

Kahneman, Daniel, *Thinking, Fast and Slow*, New York: Farrar, Straus and Giroux, 2013.

Kamenica, Emir, and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, Vol. 101, No. 6, October 2011, pp. 2590–2615.

Kennedy, Peter W., "Information Processing and Organization Design," *Journal of Economic Behavior & Organization*, Vol. 25, No. 1, 1994, pp. 37–51.

Klein, Gary, *Sources of Power: How People Make Decisions*, Cambridge, Mass.: MIT Press, 1999.

Krishna, Vijay, *Auction Theory*, 2nd ed., Burlington, Mass.: Elsevier, Academic Press, 2009.

Kydd, Andrew, "Game Theory and the Spiral Model," *World Politics*, Vol. 49, No. 3, April 1997, pp. 371–400.

Kydd, Andrew, "Which Side Are You On? Bias, Credibility, and Mediation," *American Journal of Political Science*, Vol. 47, No. 4, October 2003, pp. 597–611.

Ledvina, Andrew, and Ronnie Sircar, "Dynamic Bertrand Oligopoly," *Applied Mathematics & Optimization*, Vol. 63, No. 1, 2011, pp. 11–44.

McCarty, Nolan, and Adam Meirowitz, *Political Game Theory: An Introduction*, New York: Cambridge University Press, 2007.

McFarlan, F. Warren, "Information Technology Changes the Way You Compete," *Harvard Business Review*, Vol. 62, May 1984, pp. 98–103.

Menda, Kunal, Yi-Chun Chen, Justin Grana, James W. Bono, Brendan D. Tracey, Mykel J. Kochenderfer, and David Wolpert, "Deep Reinforcement Learning for Event-Driven Multi-Agent Decision Processes," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 20, No. 4, April 2019, pp. 1259–1268.

Miller, George A., "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information," *Psychological Review*, Vol. 63, No. 2, March 1956, pp. 81–97.

Morrow, James D., *Game Theory for Political Scientists*, Princeton, N.J.: Princeton University Press, 1994.

Mullainathan, Sendhil, and Richard H. Thaler, "Behavioral Economics," Cambridge, Mass.: National Bureau of Economic Research, NBER Working Paper 7948, October 2000.

Nash, John, "Non-Cooperative Games," *Annals of Mathematics*, Vol. 54, No. 2, September 1951, pp. 286–295.

Osborne, Martin J., *An Introduction to Game Theory*, 1st ed., New York: Oxford University Press, 2003.

Papadimitriou, Christos H., "The Complexity of Finding Nash Equilibria," in Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, eds., *Algorithmic Game Theory*, New York: Cambridge University Press, 2007, pp. 29–51.

Powell, Robert, "Bargaining and Learning While Fighting," *American Journal of Political Science*, Vol. 48, No. 2, April 2004, pp. 344–361.

Ramsay, Kristopher W., "Information, Uncertainty, and War," *Annual Review of Political Science*, Vol. 20, No. 1, 2017, pp. 505–527.

Russett, Bruce M., "Components of an Operational Theory of International Alliance Formation," *Journal of Conflict Resolution*, Vol. 12, No. 3, September 1, 1968, pp. 285–301.

Sandler, Todd, "Collective Versus Unilateral Responses to Terrorism," *Public Choice*, Vol. 124, No. 1, July 2005, pp. 75–93.

Sato, Yuzuru, Eizo Akiyama, and J. Doyne Farmer, "Chaos in Learning a Simple Two-Person Game," *Proceedings of the National Academy of Sciences*, Vol. 99, No. 7, April 2, 2002, pp. 4748–4751.

Silver, David, and Demis Hassabis, "AlphaGo: Mastering the Ancient Game of Go with Machine Learning," *Google AI Blog*, January 27, 2016.

Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al., "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature*, Vol. 529, No. 7587, January 28, 2016, pp. 484–489.

Silver, David, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al., *Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm*, ArXiv.org, December 5, 2017.

Silver, David, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al., "A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go Through Self-Play," *Science*, Vol. 362, No. 6419, December 7, 2018, pp. 1140–1144.

Simon, Herbert A., "Bounded Rationality in Social Science: Today and Tomorrow," *Mind & Society*, Vol. 1, No. 1, March 2000, pp. 25–39.

Sipser, Michael, "The History and Status of the P Versus NP Question," *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing*, New York: Association for Computing Machinery, July 1992, pp. 603–618.

Smith, Alastair, "Alliance Formation and War," *International Studies Quarterly*, Vol. 39, No. 4, December 1995, pp. 405–425.

Spence, Michael, "Job Market Signaling," *Quarterly Journal of Economics*, Vol. 87, No. 3, August 1973, pp. 355–374.

Stein, Arthur A., *Why Nations Cooperate: Circumstance and Choice in International Relations*, Ithaca, N.Y.: Cornell University Press, 1990.

Tadelis, Steven, *Game Theory: An Introduction*, Princeton, N.J.: Princeton University Press, 2013.

Tardos, Éva, and Vijay V. Vazirani, "Basic Solution Concepts and Computational Issues," in Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, eds., *Algorithmic Game Theory*, New York: Cambridge University Press, 2007, pp. 3–28.

Tromp, John, "The Number of Legal Go Positions," in Aske Plaat, Walter Kosters, and Jaap van den Herik, eds., *Computers and Games*, Cham, Switzerland: Springer International Publishing, 2016, pp. 183–190.

Tsebelis, George, "The Abuse of Probability in Political Analysis: The Robinson Crusoe Fallacy," *American Political Science Review*, Vol. 83, No. 1, March 1989, pp. 77–91.

Verbeeck, Katja, Ann Nowé, Tom Lenaerts, and Johan Parent, "Learning to Reach the Pareto Optimal Nash Equilibrium as a Team," in Bob McKay and John Slaney, eds., *AI 2002: Advances in Artificial Intelligence*, Berlin, Germany: Springer, 2002, pp. 407–418.

Welburn, Jonathan William, Justin Grana, and Karen Schwindt, "Cyber Deterrence or: How We Learned to Stop Worrying and Love the Signal," Santa Monica, Calif.: RAND Corporation, WR-1294-OSD, 2019. As of October 26, 2021:
https://www.rand.org/pubs/working_papers/WR1294.html

Wirl, Franz, "The Dynamics of Lobbying—A Differential Game," *Public Choice*, Vol. 80, No. 3, 1994, pp. 307–323.