Journal of Economic Perspectives—Volume 30, Number 4—Fall 2016—Pages 131–150

# New Directions for Modelling Strategic Behavior: Game-Theoretic Models of Communication, Coordination, and Cooperation in Economic Relationships

Vincent P. Crawford

Half a century ago, before the game-theoretic revolution that began in the 1960s and 1970s, economics largely lacked the tools to analyze strategic interactions. There was clearly a perceived need for such tools, and considerable excitement had greeted the publication of von Neumann and Morgenstern's *Theory of Games and Economic Behavior* (1944, 1947, 1953). But despite the initial excitement, for several decades game theory remained mostly a branch of mathematics, whose economic applications were the work of a few pioneers, such as Nash (1950, 1953), Schelling (1960), Shapley and Shubik (1954, 1971), and Shubik (1959). Some economists, making a virtue of presumed necessity, claimed that questions involving strategy or information were unimportant. A memorable example is Rothschild's (1973, p. 1283) quoting of a "prominent" colleague: "The friction caused by disequilibrium and lack of information accounts for variations in the numbers we observe at the fifth or sixth decimal place. Your stories are interesting but have no conceivable bearing on any question of practical economic interest."

Finally, in the 1960s, 1970s, and 1980s, game theory began to change the landscape of economics. If economists from that time could examine a modern graduate microeconomics text (such as Mas-Colell, Whinston, and Green 1995—still thoroughly "modern"), they would find their theories of market competition transformed beyond recognition, with rich, explicit game-theoretic analyses of preemption and

■ *Vincent P. Crawford is Drummond Professor of Political Economy, University of Oxford, and Fellow of All Souls College, Oxford, England. He is also Distinguished Professor Emeritus and Research Professor, University of California, San Diego, La Jolla, California. His email address is vincent.crawford@economics.ox.ac.uk.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
http://dx.doi.org/10.1257/jep.30.4.131 doi=10.1257/jep.30.4.131

entry deterrence; signalling and screening with asymmetric information; competition via explicit and/or implicit contracts; and platform and network competition. They would also find unfamiliar but flourishing subdisciplines on game-theoretic topics such as auctions; bargaining and coordination; agency and contract theory; strategic communication; social choice; public goods; cooperation in long-term relationships; and design of markets and other institutions. Such analyses, whose strategic aspects had made them seem intractable, now make up most of the microeconomics core in leading graduate programs. In the 21st century, game theory has fulfilled a large part of its promise, giving systematic, illuminating analyses of many central questions. Indeed, game theory has also begun to unify the rest of the social sciences, transforming parts of political science, computer science, and evolutionary biology—though not yet having as much effect on anthropology, sociology, or psychology.

Although most of the research that revolutionized game theory was done by economists, the revolution was not primarily a question of economics coming to game theory. Rather, game theory and economics coevolved, with game theory supplying a precise and detailed language for describing strategic interactions and a set of assumptions for predicting strategic behavior, while economics contributed questions and intuitions about behavior against which game theory's predictions could be tested and improved. In the process, the research frontier shifted from the earlier stages of figuring out how to model economic interactions as games and getting the logic of rational strategic behavior right to a later emphasis on relaxing unnecessary restrictions and refining behavioral assumptions. As game theory enriched economics, economics drove adaptations of game theory's assumptions and methods, transforming it from a branch of mathematics with a primarily normative focus into a powerful tool for positive economic analysis with a mainly descriptive or predictive focus.

In this paper, I discuss the state of progress in applications of game theory in economics and try to identify possible future developments that are likely to yield further progress. To keep the topic manageable, I focus on a canonical economic problem that is inherently game-theoretic, that of fostering efficient coordination and cooperation in relationships, with particular attention to the role of communication. I thus favor microeconomics, omitting important macroeconomic applications of game theory such as Summers (2000), Garcia-Schmidt and Woodford (2014), and Evans and McGough (2015), whose discussions of financial crises and expectations formation nonetheless touch on some of the game-theoretic issues discussed here. I also favor noncooperative game theory, omitting notable successes of cooperative game theory.[1] I further narrow the focus to problems specific to game theory

[1] The established terms "noncooperative" and "cooperative" game theory are misnomers, in that, paradoxically, noncooperative game theory is better suited to explaining (as opposed to assuming) cooperation than cooperative game theory. Noncooperative game theory starts with a detailed model of the structure of a game and makes specific assumptions about how rational players will respond to it. Cooperative game theory starts instead with a general description of the structure, sidestepping most details, and makes general assumptions intended to characterize the possible outcomes of frictionless bargaining among rational players. A notable economic application of cooperative game theory is the

by assuming that individuals are rational in the decision-theoretic sense of choosing strategies that are best responses to consistent beliefs.

I begin with an overview of noncooperative game theory's principal model of behavior, *Nash equilibrium*, henceforth shortened to *equilibrium*. I next discuss the alternative "thinking" and "learning" rationales for how real-world actors might reach equilibrium decisions. I then review how equilibrium has been used to model coordination, communication, and cooperation in relationships, and discuss possible developments. Throughout the paper, I make no attempt at comprehensive coverage or referencing, with apologies to those whose work is slighted.

## The Notion of Equilibrium in Noncooperative Game Theory

Equilibrium is defined as a combination of decision rules or *strategies*, one for each decision maker or *player*, in which each player's strategy maximizes her/his personal expected utility or *payoff* given the strategies of others who are deciding in the same way. The generality, tractability, and precision of equilibrium analysis have made it the method of choice in most economic applications of game theory (Myerson 1999). However, equilibrium goes well beyond the notion of rationality of individual decisions in that it requires a particular relationship among players' strategies. How players' strategies might come to be in equilibrium is a difficult question, which is still on the research frontier and which is intimately related to the question of how players can foster coordination and cooperation in relationships, as explained below.

Consider the game in Figure 1, in which the players choose their moves simultaneously and it is assumed that the game's structure is known to the players as *common knowledge*, in the sense that each player knows the structure, including what the other knows; knows that the other knows the structure; and so on.

This game has a unique equilibrium, in which the Row player (whose payoffs are in the lower-left corners of the cells of the matrix) chooses the strategy Middle and the Column player (whose payoffs are in the upper-right corners) chooses Center. To see this, note that if Row chooses Middle, Column will look across the choices of Left, Center, or Right, and see that Center then has the highest payoff (its two is better than the zeros for Left or Right). Further, if Column chooses Center, Row will look across the choices of Top, Middle, or Bottom, and see that Middle then has the highest payoff (its two is better than the zeros for Top or Bottom). The outcome of {Middle, Center} is therefore an equilibrium. The reader can confirm that, starting

theory of matching markets, which uses game-theoretic notions like the core to model the outcomes of competition, with or without prices, among heterogeneous traders. For more on matching theory and applications, in the context of the Nobel Memorial Prize in Economic Sciences awarded to Lloyd Shapley and Alvin Roth in 2012, see Economic Sciences Prize Committee of the Royal Swedish Academy of Sciences (2012).

*Figure 1*
**Equilibrium and Rationalizability: A Dominance-Solvable Game**

Column player

Row player

| | *Left* | *Center* | *Right* |
|---|---|---|---|
| *Top* | 0<br>7 | 5<br>0 | 3<br>0 |
| *Middle* | 0<br>5 | 2<br>2 | 0<br>5 |
| *Bottom* | 7<br>0 | 5<br>0 | 3<br>7 |

*Note:* This game has a unique equilibrium, in which the Row player (whose payoffs are in the lower-left corners of the cells of the matrix) chooses the strategy Middle and the Column player (whose payoffs are in the upper-right corners) chooses Center. A strategy choice is *strictly dominated* by another if it yields a strictly lower payoff regardless of what choice another may make. The game in Figure 1 is *dominance-solvable:* Row knows that Column is rational, and thus knows that Column will not play Right, which is strictly dominated by Center. In turn, Column knows that of the remaining choices, Row will not play Bottom, which is strictly dominated by Middle once Column's strategy Right is eliminated. Next, Row knows that of the remaining choices, Column will not play Left, which is strictly dominated by Center once Row's strategy Bottom is eliminated. The fourth step then leads precisely to the {Middle, Center} equilibrium.

from any other cell, holding one player's choice constant, the other would prefer to switch to a different choice, so {Middle, Center} is the only equilibrium.[2]

However, just knowing that {Middle, Center} is the unique equilibrium is not enough to ensure that rational players will make those choices. Suppose players have possibly probabilistic beliefs about each other's strategy choices. Then in the game in the first panel, a rational Row will play Middle only if Row's beliefs assign high enough probability to Column playing Center. Conversely, if Row's beliefs assign high probability to Column's choosing Left or Right, then Row will be tempted to play Top or Bottom. By contrast, a rational Column will never play Right, because for Column that choice is *strictly dominated*, meaning that for Column, Right yields a strictly lower payoff than Center, without regard to Row's strategy choice. But a rational Column might play Left, if Column's beliefs assign high probability to Row's choosing Bottom.

How can this ambiguity of rationality-based predictions be resolved?[3] One common approach is to strengthen the rationality assumption by making players' rationality (in addition to the structure of the game) *common knowledge*, in the sense that all players are rational, all know that all are rational, and so on ad infinitum.

[2] I ignore randomized, or *mixed*, strategies throughout the paper, and they are irrelevant to the points I make here.

[3] Manski (2003) has argued that economists should be tolerant of ambiguous predictions or as he calls them, *incomplete models*. However, his main focus is on modelling individual decisions. In games, ambiguous predictions of individual decisions frequently "multiply up" to create severe ambiguity of predicted game outcomes (Aradillas-Lopez and Tamer 2008).

*Figure 2*

**Equilibrium and Rationalizability: A Unique Equilibrium without Dominance**

Column player

| Row player | *Left* | *Center* | *Right* |
|---|---|---|---|
| *Top* | 7, 0 | 0, 5 | 0, 7 |
| *Middle* | 5, 0 | 2, 2 | 5, 0 |
| *Bottom* | 0, 7 | 0, 5 | 7, 0 |

*Note:* Like the game of Figure 1, this game also has the unique equilibrium of {Middle, Center} (Row player choses Middle and Column player chooses Center). However, this problem cannot be solved by iterated strict dominance.

Common knowledge of rationality does in fact yield a unique prediction in the game in Figure 1, which is *dominance-solvable*—meaning if players eliminate their strictly dominated strategies, and after that, their strategies that become strictly dominated once others are eliminated, and so on, the game gradually reduces to one in which only the unique equilibrium choices remain. The logic of the argument works like this: Row knows that Column is rational, and thus knows that Column will not play Right, which is strictly dominated by Center. In turn, Column knows that of the remaining choices, Row will not play Bottom, which is strictly dominated by Middle once Column's strategy Right is eliminated. Next, Row knows that of the remaining choices, Column will not play Left, which is strictly dominated by Center once Row's strategy Bottom is eliminated. The fourth step then leads precisely to the {Middle, Center} equilibrium. In dominance-solvable games whose players have more strategies, such epistemic reasoning may go on even longer before reaching equilibrium.

Now consider the game in Figure 2. It also has a unique equilibrium: If Row plays Middle, then the best choice for Column is Center; and if Column plays Center, the best choice for Row is Middle. But in that game, no choice is strictly dominated for either player, and so, even with common knowledge of rationality, epistemic logic alone does not narrow the possibilities down to a single outcome. In fact, for any strategy combination in this game, one can construct a "tower" of beliefs to show that it is consistent with common knowledge of rationality. A rational Row, for instance, might play Top because of a belief that Column will play Left (hoping for the high payoff of 7), while a rational Column might play Left because of a belief that a rational Row will play Bottom (hoping for the high payoff of 7). Some beliefs that are consistent with common knowledge of rationality lead to the equilibrium, but most do not.

More generally, Bernheim (1984) and Pearce (1984) showed that common knowledge of rationality, with no further restrictions on beliefs, implies only that each player's strategy is *rationalizable*, which can be iteratively defined as follows. A *1-rationalizable* strategy is one for which there is some profile of others' strategies that makes it a best response; a *2-rationalizable* strategy is one for which there is a profile of others' 1-rationalizable strategies that makes it a best response; and so on. A rationalizable strategy is then one that is *k*-rationalizable for all *k*. In the game in Figure 1, the choices of Middle for Row and Center for Column are both 4-rationalizable, referring to the four steps in which the players eliminate various choices via *iterated strict dominance*; and four rounds of iterated strict dominance identify the unique equilibrium. In the second game (Figure 2), all of each player's strategies are *k*-rationalizable for all *k*, and the equilibrium cannot be identified by iterated strict dominance, even though it is also unique.

The dominance-solvability of Figure 1's game is atypical in applications, as indeed is the uniqueness of the equilibrium in the games of Figure 1 and 2. As the examples suggest, equilibrium is a much stronger behavioral assumption than the rationalizability that follows from common knowledge of players' rationality. It also requires that players' beliefs be coordinated.

## Equilibrium via Thinking or Learning

If epistemic arguments based on common knowledge of rationality do not justify the coordination of players' beliefs or strategies required for equilibrium, how might it be justified? Assuming for simplicity that the structure of the game is common knowledge and that the players know that each other is rational, it is useful to sort applications into two groups according to the most plausible rationale for equilibrium: "thinking" applications, in which players can plausibly reason their way to an equilibrium; and "learning" applications, in which it is plausible that players who adjust their strategies adaptively will converge over time to some equilibrium. The rationale for assuming equilibrium affects the credibility of assuming equilibrium in applications, so I now discuss each approach.

### Thinking Applications

In thinking (as opposed to learning) applications, players play a game with no prior experience with analogous games. If assuming equilibrium is justified, it must then be because players can reason their way to equilibrium beliefs and strategy choices. In theory, this is possible if there is a commonly known principle that focuses players' beliefs on a unique prediction, because in the standard framework such common knowledge implies that their beliefs must be the same, and therefore, given rationality, in equilibrium. (For a good introduction to epistemic game theory in this journal, see Brandenberger 1992.) In this view, equilibrium becomes an *equilibrium in beliefs*, in which rational players' beliefs are statistically correct, given the best responses they imply.

Applications for which the thinking justification for equilibrium is behaviorally plausible are limited because in all but the simplest games the reasoning it requires is dauntingly complex. In Figure 1's dominance-solvable game, such reasoning requires four iterative rounds; and in the second, Figure 2, game, finding the equilibrium requires what is called "fixed-point reasoning," whereby players' strategy choices are justified as best responses to others' choices in a *two-way* recursion. (That is, one player's choice is a best response to the other's, and vice versa; dominance reasoning is also recursive, but only one-way.) In experiments that elicit subjects' initial responses to games, and that separate fixed-point and other kinds of strategic reasoning, subjects rarely follow fixed-point reasoning or indefinitely iterated dominance (Crawford, Costa-Gomes, and Iriberri 2013, Section 3). It is sometimes suggested that experienced decision makers will nonetheless use fixed-point reasoning when the stakes are high, but I have yet to find even anecdotal evidence that quants or artificial intelligence analysts of poker use reasoning that subtle.

Equilibrium reasoning becomes still more complex when the game has multiple equilibria. The logic of epistemic equilibrium-in-beliefs requires a selection among equilibria because a player who is unsure which equilibrium others have in mind will not generally find it rational to play her or his part of any particular equilibrium. Further, many important applications have multiple *strict* equilibria, in which each player has a strict preference for a strategy given others' strategies; in which case, unique equilibrium selection requires common knowledge of a complex *coordination refinement,* designed (unlike most equilibrium refinements) to discriminate among such strict equilibria. The leading examples of such refinements are from Harsanyi and Selten's (1988) classic work *A General Theory of Equilibrium Selection in Games*, which is part of the work for which they shared the 1994 Nobel prize in economics with John Nash. Their notion of *payoff-dominance* favors equilibria whose payoffs are not Pareto-inferior to those of other equilibria. Their alternative notion of *risk-dominance* favors equilibria with (roughly) larger "basins of attraction," that is, larger sets of beliefs that make their strategies best responses. Harsanyi and Selten showed that a logically consistent theory could be built on those foundations, with added tie-breaking devices, to select a unique equilibrium in any (finite matrix) game. It was a major achievement to show that such a theory could be constructed; but it rests on some unavoidably arbitrary choices and is complex enough to render it far from compelling, behaviorally (for discussion, see Aumann's "Foreword" to Harsanyi and Selten's 1988 book). But compellingness is essential in applications that involve thinking about multiple equilibria.[4]

How do people choose their strategies in thinking applications if fixed-point or indefinitely iterated dominance reasoning or equilibrium selection principles are too complex to focus their beliefs on an equilibrium? It may seem unlikely that any

[4] Some theorists believe the problem of equilibrium selection via thinking in games with multiple strict equilibria is settled by "global games" analyses (Carlsson and van Damme 1993). Such analyses add privately observed payoff perturbations to the original game in a way that makes the game dominance-solvable, and in simple coordination games makes the risk-dominant equilibrium in the unperturbed game the unique equilibrium. Although such analyses provide a systematic way to analyze how the information structure influences equilibrium selection, I believe they do not provide a conclusive argument for selecting the risk-dominant equilibrium, because the payoff perturbations are artificially introduced, and a behaviorally implausibly high number of rounds of iterated dominance are often needed to reach equilibrium in the perturbed game.

alternative model can predict observed behavior systematically better than a rational expectations notion such as equilibrium, or that such a model could be identified from among the enormous number of possible models. However, a growing body of experimental work surveyed in Crawford, Costa-Gomes, and Iriberri (2013, section 3) shows that subjects' initial responses to games often follow simple *level*-k (Costa-Gomes, Crawford, and Broseta 2001; Costa-Gomes and Crawford 2006) or *cognitive hierarchy* (Camerer, Ho, and Chong 2004) rules, in which players anchor their beliefs in a naive model of others' responses to the game and then adjust their beliefs by thinking through a small number of iterated best responses, a number which varies across players but with a stable population distribution. Such rules are decision-theoretically rational, and in sufficiently simple games they mimic equilibrium strategies. In more complex games, such rules may lead to outcomes that deviate systematically from equilibrium. Importantly, level-$k$ or cognitive hierarchy models predict not only that deviations from equilibrium will sometimes occur, but also which settings are likely to evoke them, the forms they are likely to take, and their relative frequencies. When applied to games with multiple equilibria, with estimated population frequencies of rules, they predict selection among equilibria (or not), while avoiding the complexity of coordination refinements.

The literature on strategic thinking in initial responses to games is evolving rapidly, and level-$k$ and cognitive hierarchy models are mentioned here not as the last word, but to illustrate that structural nonequilibrium models of strategic thinking are possible, and can be helpful.

**Learning Applications**

In learning applications, players have ample prior experience with closely analogous games. The learning process is modelled as repeated play of a given game, with the game that is repeated called the *stage game* and each stage game normally with a different partner.[5] Players' choices are modelled as *adaptive learning*, in which they adjust their stage-game strategies over time in ways that increase their own stage-game payoffs on the (usually false) assumption that others' stage-game strategies will continue as before. In adaptive learning models, players have a strong tendency to converge to some equilibrium in the stage game. There are few general theoretical results, but there is strong experimental support for such convergence.

Even if learning assures convergence to some equilibrium, nonequilibrium strategic thinking often remains relevant. Suppose that only long-run outcomes matter, but the stage game in the applications has multiple equilibria, as in many important applications. Then all we need from game theory is a reliable prediction of the prior probability distribution of the possible equilibrium outcomes. But with

[5] Unless players' partners vary across the stage games, repeated-game strategies are relevant, and it is implausible that players focus on their choices of stage-game strategies, stage by stage, as opposed to thinking about the effectiveness of their alternative repeated-game strategies. I ignore the literature on *rational learning* models, in which players are assumed to play an equilibrium in the repeated game that describes the entire learning process, because this approach seems less useful than adaptive learning models in applications (for example, Crawford 2001, Section 6.4.4).

multiple equilibria, learning dynamics are normally history-dependent, so people's initial responses influence limiting outcomes, as do the structures of their learning rules (Van Huyck, Cook, and Battalio 1997; Crawford 1995, 2001; Camerer and Ho 1998, 1999).[6] Moreover, even if the stage game has a unique equilibrium, analogies across games are rarely as close in applications as adaptive learning models assume, and analysis of the thinking needed to learn from imperfect analogies has just begun (Rankin, Van Huyck, and Battalio 2000; Van Huyck and Battalio 2002; Cooper and Kagel 2003; Samuelson 2001).

## Communication, Coordination, and Cooperation in Economic Relationships

To perform well, an economic relationship must solve one or more strategic problems. For example, players may face *incentive problems* that encourage them not to cooperate, even though cooperating would be in both of their interests, as in the well-known Prisoner's Dilemma game discussed further below. Players may face *assurance problems* that make it seem too risky to play the strategies that would lead to efficient equilibria, as in the Stag Hunt game discussed below. Players may also face *bargaining/coordination problems* that make it difficult to coordinate on one of multiple efficient equilibria.[7] To be useful in applications, game theory must offer theoretically coherent and behaviorally credible analyses of these kinds of problems.

In the standard approach to these problems, in one-shot interactions all three problems are solved (or not) within an equilibrium, sometimes augmented by principles that govern selection among multiple equilibria. In the standard approach to repeated interactions, all three problems are again solved (or not) within an equilibrium, now with refinements like *subgame-perfect equilibrium* applied to the game that describes the entire relationship.[8] Solving strategic problems may look quite different in situations of static play or repeated play—although if a game is repeated with fixed rather than varying partners, it becomes a game that is effectively static,

[6] Some theorists consider the problem of equilibrium selection via learning to be settled by analyses of "long-run equilibria" (Kandori, Mailath, and Rob 1993; Young 1993). But those analyses achieve equilibrium selection by modelling the dynamics of learning as ergodic and passing to the limit as randomness in the dynamics becomes negligible. Neither feature seems realistic, nor do the results seem to correspond closely to equilibrium selection in the lab or the field (Crawford 2001).

[7] Here I follow Schelling (1960) and Roth (1987; see also Crawford 1997, Section 5.3) in suggesting that most real bargaining is best modelled as unstructured and is then primarily a coordination problem, not a problem that is resolved via delay costs in the subgame-perfect equilibrium of a game with artificially imposed timing of offers and counteroffers (Rubinstein 1982).

[8] A *subgame* is any part of a game that remains after part of it has been played. A *subgame-perfect equilibrium* is an equilibrium strategy profile that induces an equilibrium in every subgame. In effect, subgame-perfect equilibrium adds a time-consistency requirement to the notion of equilibrium. This notion can be generalized to games with asymmetric information, via notions called "sequential equilibrium" or "perfect Bayesian equilibrium" (Mas-Colell, Whinston, and Green 1995, chapters 8–9).

*Figure 3*
**Prisoner's Dilemma**

| | *Cooperate* | *Defect* |
|---|---|---|
| *Cooperate* | 3 / 3 | 5 / 0 |
| *Defect* | 0 / 5 | 1 / 1 |

with players making a one-time choice among strategies that describe how they will act as the game unfolds.

If players can communicate during their interactions, that is usually modelled via "cheap talk" messages, which involve no direct payoff consequences and have no power to commit players to actions (Crawford and Sobel 1982; or in this journal, Farrell and Rabin 1996).

In this section, I begin to explore new directions for modelling communication, coordination, and cooperation in relationships. I first explain the standard approaches, and then suggest alternative directions that seem likely to be feasible and potentially useful.

### Coordination and Cooperation in Long-Term Relationships

Most existing work in repeated games seeks to identify ways to support cooperation in a subgame-perfect equilibrium of the infinitely repeated Prisoner's Dilemma or some other well-behaved repeated game (Fudenberg and Maskin 1986). Consider, for instance, the version of the Prisoner's Dilemma in Figure 3. The best symmetric outcome arises if both players choose Cooperate. However, Defect is a dominant strategy for each player in that Defect yields each player a strictly higher payoff than Cooperate whether or not the other player chooses Cooperate. Is there a way to support cooperation in equilibrium in relationships based on the Prisoner's Dilemma?

First suppose that the Prisoner's Dilemma in Figure 3 is played repeatedly for a potentially infinite number of times by the same two players; and after any given number of plays the conditional probability of continuing remains bounded above zero, so it never becomes common knowledge that any particular period will be the last one in the relationship. Assume that players will choose a strategy to maximize their payoffs added across plays of the game, but downweight future payoffs with a discount factor (because otherwise undiscounted payoffs over an infinite horizon are not well-defined). A relatively low discount factor means that the players do not place much weight on future payoffs. In that case, only repeated choices of {Defect, Defect} are consistent with equilibrium. By contrast, a high discount factor means that the players place high weight on future payoffs; and that is enough to make {Cooperate, Cooperate} consistent with subgame-perfect equilibrium. For example, both players could follow the "grim trigger" strategy "Cooperate until the other player Defects, then Defect forever," which happens to be a subgame-perfect equilibrium and yields the outcome {Cooperate, Cooperate} in every period.

Why is it important that the game be played a potentially infinite number of times? Suppose instead that the Prisoner's Dilemma in Figure 3 is played repeatedly for a commonly known, finite number of times by the same two players. Assume again that players' preferences are defined by the addition of players' payoffs across plays of the game, with or without discounting. The unique equilibrium then entails both players choosing Defect in every period. During the last period of the game—which is known in advance—Defect is a dominant strategy for both players. Knowing that, in the second-to-last period, the players recognize that they cannot avoid Defect-Defect in the last period, and Defect is therefore a conditionally dominant strategy for both. Working backward, both players will see that Defect is a dominant strategy in every period.

Most existing work in repeated games assumes that players' beliefs will focus, with certainty, on a particular subgame-perfect equilibrium as common knowledge, and seeks to characterize the "Folk Theorem" set of outcomes, those which are consistent with some such subgame-perfect equilibrium (for more on the Folk Theorem, see the Wikipedia entry: https://en.wikipedia.org/wiki/Folk_theorem_(game_theory)). However, with a potentially infinite horizon, the set of equilibria is usually enormous. In the Prisoner's Dilemma in Figure 3, consider the strategy combination: "Row initially Cooperates and then alternates between Defect and Cooperate, and Column always Cooperates—in each case until either player deviates (that is, Row Defecting when not supposed to, or Column Defecting at all), in which case both players Defect from now on." In this asymmetric strategy combination, Row does better and Column worse than in the symmetric equilibrium described above (where both use the grim-trigger strategy), but if the discount factor is high enough, then the punishments for deviations are costly enough to make it a subgame-perfect equilibrium as well. And there are many others.

This multiplicity of equilibria is an important difficulty because, in applications, uncertainty about one's partner's strategic thinking is of the essence. The complexity of repeated-game equilibria and the general difficulty of equilibrium selection make the thinking justification especially implausible here. And in real long-term relationships, players' opportunities for learning about the effectiveness of alternative repeated-game strategies are limited (but see Dal Bó and Fréchette 2011 for some intriguing experimental evidence on learning repeated-game strategies).

Applications of such repeated-games analyses must confront a number of issues, of which I mention four. First, Folk Theorem equilibria are normally supported by extreme punishments even for tiny deviations. (The punishments are often taken to be more extreme than necessary to support cooperation because that gives a cleaner characterization of the Folk Theorem set.) Imagine a relationship slightly more complicated than a repeated Prisoners' Dilemma, whose players start out with different beliefs about their repeated-game strategies: For example, one player might believe they are playing "Cooperate until a player defects, then defect forever," while the other believes they are playing the asymmetric strategy combination described in the previous paragraph. In this case, they will deviate while intending to cooperate, and the trigger strategies meant to support their cooperation will end cooperation. This brittleness suggests that in applications, people will favor strategies that are more robust to deviations. There are few such analyses, but see Porter (1983), van Damme (1989), and Friedman and Samuelson (1994).

A second issue involves the ambiguity of predictions associated with the extreme multiplicity of equilibria in repeated games analyses. This ambiguity has been a serious impediment to empirical applications, and I believe that it has slowed the co-evolution of theory, experiment, and empirics that has been such a powerful engine of progress in other parts of game theory. Perhaps surprisingly, there seems reason to hope that closer attention in theoretical analyses to the need for strategies to be robust will, as a side benefit, help reduce ambiguity of predictions of players' behavior. Recent experimental work by Blonski, Ockenfels, and Spagnolo (2011), Breitmoser (2015), and others suggests the possibility of better and more precise theory.

A third issue is that long-term relationships enable *strategic teaching*, in which a player whose future cooperation with current partners is worth preserving may try to benefit by deviating from a short-run payoff-maximizing strategy in a way that could influence others' future beliefs and choices (Camerer, Ho, and Chong 2002). For instance, in the repeated Prisoner's Dilemma game of Figure 3, Row would benefit if it were possible to teach Column to play the asymmetric equilibrium described above (Row initially Cooperates and then alternates between Defect and Cooperate, and Column always Cooperates—in each case until either player deviates) rather than the symmetric equilibrium in which both players follow the "grim trigger" strategy, which shares the surplus equally. Row could try to teach Column by deviating from the latter equilibrium, risky as that is. Such considerations highlight the importance of robustness, but are assumed away in a standard equilibrium analysis.

Van Huyck, Battalio, and Beil's (1990) experiments with two-person minimum-effort coordination games, like the Stag Hunt game presented below but with seven symmetric Pareto-ranked equilibria, provide an intriguing example of strategic teaching. When their subjects played the games in fixed pairs, but with only one repetition per play, many of them adjusted their current decisions to try to teach their partners to coordinate more efficiently, and 12 of the 14 subject pairs converged via various routes to the most efficient equilibrium. (By contrast, subjects in the analogous treatment with random re-pairing of partners did not try to teach their partners, and had significantly worse outcomes.) The puzzle is how did they learn enough about the effectiveness of their alternative repeated game strategies to play the efficient stage-game equilibrium? In Crawford (2002), I suggested that Van Huyck et al.'s results might be explained by a strategic teaching model like that of Camerer, Ho, and Chong (2002), in which some players are adaptive learners while others are forward-looking and *sophisticated* in the sense of best responding to the correct mixture of adaptive and sophisticated subjects.

The fourth and last issue I will mention here is that in most standard repeated-game models, players who are well-informed about the structure of the game have nothing to communicate in equilibrium, so such models imply no substantive role

*Figure 4*
**Stag Hunt**

| | *Stag* | *Rabbit* |
|---|---|---|
| *Stag* | 9, 9 | 1, 8 |
| *Rabbit* | 8, 1 | 7, 7 |

for communication (for a recent exception, see Awaya and Krishna 2016). Yet communication appears to interact in important ways with the phenomena just discussed, and to play an essential role in real relationships.

**Using Communication to Foster Coordination and Cooperation**

Humans appear to be uniquely capable of using language to build, communicate, and counterfactually manipulate mental models of the world and of other people. This capability has a powerful influence on how people structure and maintain their relationships, and on what they can accomplish in them. Yet existing models of collusion and cooperation assign a limited role to communication. For example, most repeated-games analyses imply that firms can accomplish as much via tacit collusion as with communication. Why, then, does American antitrust law bother to prohibit firms from communicating about pricing and output decisions (Genesove and Mullin 2001; Andersson and Wengström 2007)? Presumably, when an agreement has gone awry, communication is an important aid to understanding what went wrong and restoring the relationship. Better models of how communication helps are needed.

To begin to explore these issues, consider the well-known Stag Hunt game, which traces back to a scenario laid out by Rousseau (1754 [1973]), who wrote:

> If a deer was to be taken, everyone saw that, in order to succeed, he must abide faithfully by his post: but if a hare happened to come within the reach of any one of them, it is not to be doubted that he pursued it without scruple, and, having seized his prey, cared very little, if by so doing he caused his companions to miss theirs.

A two-player Stag Hunt game with a set of payoffs is shown in Figure 4. The game has two pure-strategy equilibria, "all-Stag" and "all-Rabbit." All-Stag is better for both players than all-Rabbit, and is therefore "payoff-dominant" and a preferable equilibrium using one of the Harsanyi and Selten (1987) criteria. But as Rousseau's scenario suggests, how can the two players build trust that they will stay at their posts so that each can get a stag, rather having one of them deviate and try to bag a Rabbit? Rabbit also has a fairly large payoff, and there are far larger sets of players' beliefs that make Rabbit a best response. For the payoffs given in Figure 4, a player finds it optimal to play Rabbit if the belief is that the player's partner will play Rabbit with probability at least 1/7, while it is optimal to play Stag only under the belief that partner will play Stag with probability at least 6/7. Thus, using another of the Harsanyi and Selton criteria for choosing between equilibria, the all-Rabbit equibrium is "risk-dominant."

Experiments suggest that if people play Stag Hunt with no opportunity to communicate, a large majority of them will play Rabbit, as in other settings with a strongly risk-dominant equilibrium (Straub 1995). But now imagine, following Aumann (1990; see also Farrell 1988), that Stag Hunt is to be played only once, but that before play, one player, the *sender*, must send a clear message about the sender's intended strategy, Stag or Rabbit. As already noted, in game theory such communication is usually modelled via cheap talk messages, which are nonbinding and have no direct payoff consequences. Even so, such a message might benefit the sender by influencing the *receiver*'s choice (Crawford and Sobel 1982; Farrell and Rabin 1996).

Aumann (1990) notes that whether or not the sender plans to play Stag, the sender prefers that the receiver play Stag ($9 > 1$ and $8 > 7$). He argues that for this reason, the receiver will infer that the sender's message is self-interested and the message can convey no information to the receiver, so that the outcome will be the same as without communication. Aumann's argument is related to Farrell and Rabin's (1996) distinction between messages that are "self-committing" in that if the message convinces the receiver, it's a best response for the sender to do as he said; and those that are "self-signaling" in that they are sent when and only when the sender intends to do as he said. In this case, a message of intention to play Stag is self-committing, but not self-signaling. Aumann's argument is correct as a matter of logic, yet many of us would expect most senders to send *and* play Stag, and most receivers to play Stag as well. This conclusion is confirmed in most experiments (Cooper, DeJong, Forsythe, and Ross 1992; Charness 2000; Ellingsen, Östling, and Wengström 2013; but see Clark, Kay, and Sefton 2001).

One reason for the discrepancy has to do with Aumann's (1990) exclusive reliance on the logic of equilibrium, even though the multiplicity of equilibria, with one payoff-dominant and another risk-dominant, seriously undermines the thinking justification for equilibrium. When uncertainty about others' thinking is of the essence, it is unlikely that intelligent people will interpret a sender's message as if there were no chance whatsoever that it would influence equilibrium selection or whether players' choices are even in equilibrium. Rabin (1994; see also Farrell 1987, 1988) relaxes the assumption that players' beliefs are perfectly coordinated on some equilibrium, using a combination of rationalizability and behaviorally plausible assumptions about how players use language to analyze the process of negotiating how to play one of a class of finite matrix games. He shows that if players can communicate as long as desired, they will use their messages to agree on an equilibrium that is no worse for either player than the worst Pareto-efficient equilibrium for that player—thus, for example, yielding all-Stag in Stag Hunt.

Ellingsen and Östling (2010) take a different nonequilibrium approach, adapting the level-*k* model of Crawford (2003) to resolve some puzzles regarding the comparative effectiveness of one- or two-sided communication in coordination and other games.[9] They show, among other things, that even one-sided communication may allow players to coordinate on a Pareto-dominant equilibrium in a wide class of games including Stag Hunt, again resolving the puzzle. Notable recent experimental work includes Andersson and Wengström (2012) and Cooper and Kühn (2014), who study communication and renegotiation in two-stage games.

A more subtle reason for the discrepancy between Aumann's (1990) prediction and prevailing intuitions and experimental evidence on the effectiveness of communication in Stag Hunt may be his assumption that players are limited to a fixed list of messages, as in most theoretical work on communication: in his case, strategy labels whose meanings are assumed to be understood. Yet Stag Hunt is one of many situations in which people, even if well-informed about the structure, might benefit from a discussion more nuanced than stating an intention before deciding how to play. A sender who could send an unrestricted natural-language message would probably try to convey not only an intention but also a broader understanding of strategic issues. A fuller message might say, trying to give the assurance needed to support All-Stag: "I can see, as I am sure you can, that the best outcome in this game would be for both of us to play Stag. But I realize that Stag is risky for you, as it is for me. Despite the risk, I have concluded that Stag is a better bet for us. I plan to play Stag, and I hope you will too." In experiments such natural-language messages can be very effective (Charness and Dufwenberg 2006, 2010).

How could game theory incorporate such richer communication? Relaxing the standard assumption that people are limited to a fixed list of messages about intentions or private information to allow "metatheoretical" messages like my quotation is a theoretical challenge, and it seems difficult even to formalize an epistemic thinking justification for equilibrium when there is uncertainty about the principles of equilibrium selection, or even whether such principles ensure that players play some equilibrium. Even so, the gains from understanding natural-language communication, and how it interacts with people's other decisions, seem likely to be very large. McGinn, Thompson, and Bazerman (2003) report experimental evidence on how subjects use natural-language messages, which may help in devising better theories (on this point, see also Valley, Thompson, Gibbons, and Bazerman 2002; Weber and Camerer 2003; Charness and Dufwenberg 2006; Houser and Xiao 2011; Burchardi and Penczynski 2014).

Better theories of communication will certainly include some elements of players' rationality and their knowledge of the rationality of others, but they cannot be entirely epistemic. Rather, such theories are likely to combine rationality-based

[9] In Crawford (2003), I studied deceptive preplay communication of intentions before a zero-sum two-person game, which can happen in a plausible level-*k* model, but not in equilibrium. This level-*k* model has a more plausible thinking justification than equilibrium and also has some experimental support (Wang, Spezio, and Camerer 2010).

reasoning about the meaning of messages with empirically-based restrictions, such as those used by Farrell (1987) and Rabin (1994) and, or perhaps like those embodied in level-*k* rules like those studied by Ellingsen and Östling (2010).

## Conclusion

Some economists seem less excited about game theory than during the period in the 1960s, 1970s, and 1980s when the ability to analyze strategic interactions was altering the landscape of many subfields of economics. But if the excitement over game theory has in fact diminished, I do not believe it is because game theory has ceased to be a major driving force in economics—quite the contrary!—it is mainly because its centrality makes economists less aware of its presence. Modern economists' relationship to game theory may resemble fully adapted aquatic creatures' relationship to water: they are less aware of water than their amphibian ancestors, for whom swimming was always a choice, but also more agile in their new medium.

That said, if game theory is to continue as a major force for progress in economics, it must continue to co-evolve with economic applications and incorporate the empirical knowledge they provide, rather than pursuing an inwardly focused agenda. In this paper, I have tried to give some concrete illustrations of what that might mean, critiquing existing game-theoretic approaches to the canonical problem of using communication to foster coordination and cooperation in relationships and suggesting some directions in which further progress might be made.

■ *I thank Colin Camerer, Gary Charness, Miguel Costa-Gomes, Martin Dufwenberg, Nagore Iriberri, Robert Östling, H. Peyton Young, the editors, and especially Joel Sobel and Tore Ellingsen for their helpful discussions and comments. My research received primary funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. 339179. The contents reflect only my views and not the views of the ERC or the European Commission, and the European Union is not liable for any use that may be made of the information contained therein. I also thank All Souls College, Oxford, and the University of California, San Diego, for research support.*

## References


**Andersson, Ola, and Erik Wengström.** 2007. "Do Antitrust Laws Facilitate Collusion? Experimental Evidence on Costly Communication in Duopolies." *Scandinavian Journal of Economics* 109(2): 321–39.

**Andersson, Ola, and Erik Wengström.** 2012. "Credible Communication and Cooperation: Experimental Evidence from Multi-Stage Games." *Journal of Economic Behavior and Organization* 81(1): 207– 219.

**Aradillas-Lopez, Andres, and Elie Tamer.** 2008. "The Identification Power of Equilibrium in Simple Games." *Journal of Business & Economic Statistics* 26(3): 261–83.

**Aumann, Robert.** 1990. "Nash-Equilibria Are Not Self-Enforcing." In *Economic Decision-Making: Games, Econometrics and Optimisation*, edited by J. J. Gabszewicz, J.-F. Richard, and L. Wolsey, 201–206. North Holland.

**Awaya, Yu, and Vijay Krishna.** 2016. "On Communication and Collusion." *American Economic Review* 106(2): 285–315.

**Bernheim, B. Douglas.** 1984. "Rationalizable Strategic Behavior." *Econometrica* 52(4): 1007–1028.

**Binmore, Ken, John McCarthy, Giovanni Ponti, Larry Samuelson, and Avner Shaked.** "A Backward Induction Experiment." *Journal of Economic Theory* 104(1): 48–88.

**Blonski, Matthias, Peter Ockenfels, and Giancarlo Spagnolo.** 2011. "Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence." *American Economic Journal: Microeconomics* 3(3): 164–92.

**Brandenburger, Adam.** 1992. "Knowledge and Equilibrium in Games." *Journal of Economic Perspectives* 6(4): 83–101.

**Breitmoser, Yves.** 2015. "Cooperation, But No Reciprocity: Individual Strategies in the Repeated Prisoner's Dilemma." *American Economic Review* 105(9): 2882–2910.

**Burchardi, Konrad B., and Stefan P. Penczynski.** 2014. "Out of Your Mind: Eliciting Individual Reasoning in One Shot Games." *Games and Economic Behavior* 84(1): 39–57.

**Camerer, Colin, and Teck-Hua Ho.** 1998. "Experience-Weighted Attraction Learning in Coordination Games: Probability Rules, Heterogeneity, and Time-Variation." *Journal of Mathematical Psychology* 42(2–3): 305–326.

**Camerer, Colin, and Teck-Hua Ho.** 1999. "Experience-Weighted Attraction Learning in Normal Form Games." *Econometrica* 67(4): 837–74.

**Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong.** 2002. "Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games." *Journal of Economic Theory* 104(1): 137–88.

**Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong.** 2004. "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics* 119(3): 861–98.

**Carlsson, Hans, and Eric van Damme.** 1993. "Global Games and Equilibrium Selection." *Econometrica* 61(5): 989–1018.

**Charness, Gary.** 2000. "Self-Serving Cheap Talk: A Test of Aumann's Conjecture." *Games and Economic Behavior* 33(2): 177–94.

**Charness, Gary, and Martin Dufwenberg.** 2006. "Promises and Partnership." *Econometrica* 74(6): 1579–1601.

**Charness, Gary, and Martin Dufwenberg.** 2010. "Bare Promises: An Experiment." *Economics Letters* 107(2): 281–83.

**Clark, Kenneth, Stephen Kay, and Martin Sefton.** 2001. "When Are Nash Equilibria Self-Enforcing? An Experimental Analysis." *International Journal of Game Theory* 29(4): 495–515.

**Cooper, David J., and John H. Kagel.** 2003. "Lessons Learned: Generalizing Learning Across Games." *American Economic Review* 93(2): 202–207.

**Cooper, David J., and Kai-Uwe Kühn.** 2014. "Communication, Renegotiation, and the Scope for Collusion." *American Economic Journal: Microeconomics* 6(2): 247–78.

**Cooper, Russell, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross.** 1992. "Communication in Coordination Games." *Quarterly Journal of Economics* 107(2): 739–771.

**Costa-Gomes, Miguel A., and Vincent P. Crawford.** 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review* 96(5): 1737–68.

**Crawford, Vincent P.** 1995. "Adaptive Dynamics in Coordination Games." *Econometrica* 63(1): 103–143.

**Crawford, Vincent P.** 1997. "Theory and Experiment in the Analysis of Strategic Interaction." Chap. 7 in *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress,* vol. 1, edited by David Kreps and Kenneth F. Wallis. Cambridge University Press (Reprinted in 2003 as chap. 12 in *Advances in Behavioral Economics*, edited by Colin F. Camerer, George Loewenstein, and Matthew Rabin, Princeton University Press).

**Crawford, Vincent P.** 2001. "Learning Dynamics, Lock-in, and Equilibrium Selection in Experimental Coordination Games." Chap. 6 in *The Evolution of Economic Diversity,* edited by Ugo Pagano and Antonio Nicita. Routledge. (With correction of the published version of Figure 6.2(b) (2b in the working paper) at http://econweb.ucsd.edu/%7Evcrawfor/9719.pdf.)

**Crawford, Vincent P.** 2002. "Introduction to Experimental Game Theory." *Journal of Economic Theory* 104(1): 1–15.

**Crawford, Vincent P.** 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." *American Economic Review* 93(1): 133–49.

**Miguel Costa-Gomes, Crawford, Vincent P., and Bruno Broseta.** 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica* 69(5): 1193–1235.

**Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri.** 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." *Journal of Economic Literature* 51(1): 5–62.

**Crawford, Vincent P., and Joel Sobel.** 1982. "Strategic Information Transmission." *Econometrica* 50(6): 1431–51.

**Dal Bó, Pedro, and Guillaume R. Fréchette.** 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review* 101(1): 411–429.

**Economic Sciences Prize Committee of the Royal Swedish Academy of Sciences.** 2012. "Stable Allocations and the Practice of Market Design." Scientific Background on the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2012. http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2012/advanced-economicsciences2012.pdf.

**Ellingsen, Tore, and Robert Östling.** 2010. "When Does Communication Improve Coordination." *American Economic Review* 100(4): 1695–1724.

**Ellingsen, Tore, Robert Östling, and Erik Wengström.** 2013. "How Does Communication Affect Beliefs?" June 13. http://perseus.iies.su.se/~rob/papers/Ellingsen_et_al_2013.pdf.

**Evans, George W., and Bruce McGough.** 2015. "The Neo-Fisherian View and the Macro Learning Approach." http://economistsview.typepad.com/economistsview/2015/12/the-neo-fisherian-view-and-the-macro-learning-approach.html.

**Farrell, Joseph.** 1987. "Cheap Talk, Coordination, and Entry." *RAND Journal of Economics* 18(1): 34–39.

**Farrell, Joseph.** 1988. "Communication, Coordination, and Nash Equilibrium." *Economics Letters* 27(3): 209–214.

**Farrell, Joseph, and Matthew Rabin.** 1996. "Cheap Talk." *Journal of Economic Perspectives* 10(3): 103–118.

**Friedman, James W., and Larry Samuelson.** 1994. "The 'Folk Theorem' for Repeated Games and Continuous Decision Rules." Chap. 6 in *Problems of Coordination in Economic Activity*, edited by James W. Friedman, in Recent Economic Thought Series, vol. 35. Springer Verlag.

**Fudenberg, Drew, and Eric Maskin.** 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* 54(3): 533–554.

**García-Schmidt, Mariana, and Michael Woodford.** 2014. "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." http://www.columbia.edu/~mw2230/GSW.pdf.

**Genesove, David, and Wallace P. Mullin.** 2001. "Rules, Communication, and Collusion: Narrative Evidence from the Sugar Institute Case." *American Economic Review* 91(3): 379–98.

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

**Houser, Daniel, and Erte Xiao.** 2011. "Classification of Natural Language Messages Using a Coordination Game." *Experimental Economics* 14(1): 1–14.

**Kandori, Michihiro, George J. Mailath, and Rafael Rob.** 1993. "Learning, Mutation, and Long Run Equilibria in Games." *Econometrica* 61(1): 29–56.

**Manski, Charles F.** 2003. *Partial Identification of Probability Distributions.* Springer-Verlag.

**Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green.** 1995. *Microeconomic Theory.* Oxford University Press.

**McGinn, Kathleen L., Leigh Thompson, and Max H. Bazerman.** 2003. "Dyadic Processes of Disclosure and Reciprocity in Bargaining with Communication." *Journal of Behavioral Decision Making* 16(1): 17–34.

**Myerson, Roger B.** 1999. "Nash Equilibrium and the History of Economic Theory." *Journal of Economic Literature* 37(3): 1067–82.

**Nash, John.** 1950. "The Bargaining Problem." *Econometrica* 18(2): 155–162.

**Nash, John.** 1953. "Two-Person Cooperative Games." *Econometrica* 21(1): 128–40.

**Pearce, David G.** 1984. "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica* 52(4): 1029–50.

**Porter, Robert H.** 1983. "Optimal Cartel Trigger Price Strategies." *Journal of Economic Theory* 29(2): 313–38.

**Rabin, Matthew.** 1994. "A Model of Pre-Game Communication." *Journal of Economic Theory* 63(2): 370–91.

**Rankin, Frederick, John Van Huyck, and Ray Battalio.** 2000. "Strategic Similarity and Emergent Conventions: Evidence from Similar Stag Hunt

Games." *Games and Economic Behavior* 32(2): 315–37.

**Roth, Alvin E.** 1987. "Bargaining Phenomena and Bargaining Theory." Chap. 2 in *Laboratory Experimentation in Economics: Six Points of View*, edited by Alvin E. Roth. Cambridge University Press.

**Rothschild, Michael.** 1973. "Models of Market Organization with Imperfect Information: A Survey." *Journal of Political Economy* 81(6): 1283–1308.

**Rousseau, Jean-Jacques.** 1754 [1973]. "A Discourse on the Origin of Inequality." In *The Social Contract and Discourses* (translated by G. D. H. Cole), 27–113. London: J. M. Dent & Sons, Ltd. Also at http://www.constitution.org/jjr/ineq.htm.

**Rubinstein, Ariel.** 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50(1): 97–109.

**Samuelson, Larry.** 2001. "Analogies, Adaptation, and Anomalies." *Journal of Economic Theory* 97(2): 320–66.

**Schelling, Thomas C.** 1960. *The Strategy of Conflict*. Harvard University Press.

**Shapley, Lloyd S., and Martin Shubik.** 1954. "A Method for Evaluating the Distribution of Power in a Committee System." *American Political Science Review* 48(3): 787–92.

**Shapley, Lloyd S., and Martin Shubik.** 1971. "The Assignment Game I: The Core." *International Journal of Game Theory* 1(1): 111–30.

**Shubik, Martin.** 1959. *Strategy and Market Structure: Competition, Oligopoly, and the Theory of Games*. John Wiley and Sons, Inc.

**Straub, Paul G.** 1995. "Risk Dominance and Coordination Failures in Static Games." *Quarterly Review of Economics and Finance* 35(4): 339–63.

**Summers, Lawrence H.** 2000. "International Financial Crises: Causes, Prevention, and Cures." *American Economic Review* 90(2): 1–16.

**Valley, Kathleen, Leigh Thompson, Robert Gibbons, and Max H. Bazerman.** 2002. "How Communication Improves Efficiency in Bargaining Games." *Games and Economic Behavior* 38(1): 127–55.

**van Damme, Eric.** 1989. "Stable Equilibria and Forward Induction." *Journal of Economic Theory* 48(2): 476–96.

**Van Huyck, John, and Raymond Battalio.** 2002. "Prudence, Justice, Benevolence, and Sex: Evidence from Similar Bargaining Games." *Journal of Economic Theory* 104(1): 227–46.

**Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil.** 1990. "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review* 80(1): 234–48.

**Van Huyck, John B., Joseph Cook P., and Raymond C. Battalio.** 1997. "Adaptive Behavior and Coordination Failure." *Journal of Economic Behavior and Organization* 32(4): 483–503.

**von Neumann, John, and Oskar Morgenstern.** 1944, 1947, 1953. *Theory of Games and Economic Behavior*. Princeton University Press.

**Wang, Joseph Tao-yi, Michael Spezio, and Colin F. Camerer.** 2010. "Pinocchio's Pupil: Using Eyetracking and Pupil Dilation to Understand Truth-telling and Deception in Sender-Receiver Games." *American Economic Review* 100(3): 984–1007.

**Weber, Roberto A., and Colin F. Camerer.** 2003. "Cultural Conflict and Merger Failure: An Experimental Approach." *Management Science* 49(4): 400–415.

**Young, H. Peyton.** 1993. "The Evolution of Conventions." *Econometrica* 61(1): 57–84.