

Stochastic game theory: for playing games, not just for doing theory

Jacob K. Goeree and Charles A. Holt*

Department of Economics, Rouss Hall, University of Virginia, Charlottesville, VA 22903

Recent theoretical advances have dramatically increased the relevance of game theory for predicting human behavior in interactive situations. By relaxing the classical assumptions of perfect rationality and perfect foresight, we obtain much improved explanations of (i) initial decisions, (ii) dynamic patterns of learning and adjustment, and (iii) equilibrium steady-state distributions.

Introduction

About fifty years ago, John Nash walked into the office of the Chair of the Princeton Mathematics Department with a solution concept for N -person games, along with an existence proof that was soon published in these *Proceedings* [1]. John von Neumann was dismissive and remarked "that's trivial, you know. That's just a fixed point theorem" [2]. But word of Nash's theorem spread quickly at RAND on the West Coast, where researchers working on defense strategy were dissatisfied with the received theory of zero-sum games, since the assumption that one player's gain is another's loss is of limited relevance beyond simple card games. Two mathematicians, Dresher and Flood, designed a laboratory experiment to test Nash's equilibrium concept the same day they heard about his proof. Their experiment implemented a game in which two players have unilateral incentives to "defect" even though both are better off when both "cooperate." Nash's thesis advisor, Tucker, later saw the payoffs for this experiment on the blackboard in someone's office and devised the famous story of the "prisoner's dilemma," which he used in a seminar for the Psychology Department at Stanford University [3].

The applications of game theory have expanded greatly since then, and with the Nash equilibrium as its centerpiece, game theory has finally gained the central role first envisioned by von Neumann and Morgenstern [4]. If anything, game theory is the leading contender for

* Reprint requests should be sent by email to Goeree at jg2n@virginia.edu. This research was funded in part by the National Science Foundation (SBR-9617784 and SBR-9818683). We wish to thank Melayne McInnes for a helpful suggestion.

becoming a general theory of social science, with extensive applications in economics, political science, psychology, law, and biology. Indeed, in some areas of economics virtually all recent theoretical developments are applications of game theory.

There is, however, widespread criticism of theories based on the classical "rational choice" assumptions of perfect decision making (no errors) and perfect foresight (no surprises), especially when they are applied to describe behavior in complex interactive situations. This skepticism is reinforced by evidence from laboratory experiments with financially-motivated subjects which often produce behavior patterns that are systematically biased away from rational choice predictions. Nash himself participated in experiments as a subject and later designed experiments of his own, but he and his coauthors lost whatever confidence they had in game theory when they saw how poorly it predicted actual behavior [2]. And Reinhard Selten, who shared the 1995 Economics Nobel Prize (with Nash and Harsanyi), remarked that "game theory is for doing theory, not for playing games." Like many others, he has argued that decisions are stochastic or "noisy," where the noise in subjects' behavior may be due to errors in perception, calculation, or recording decisions [5]. Alternatively, apparent noise may represent fully rational responses to factors like benevolence, envy, or other idiosyncratic factors that are not measured by the experimenter [6]. Regardless of the source and interpretation of the noise, the effect will be that different players encounter different histories of others' play, and learning in such environments may lead to variations in individuals' beliefs and decisions. This paper describes three new developments in game theory that relax the classical assumptions of perfect rationality and perfect foresight. These approaches to noisy introspection (prior to play), learning (from previous plays), and equilibrium (after a large number of plays) provide complementary perspectives for explaining actual behavior in a wide variety of games.

Coordination and Social Dilemma Games

The models summarized here have been strongly influenced by data from experiments that show disturbing differences between game-theoretic predictions and behavior of human subjects who are earning money in controlled strategic situations. For example, Goeree and Holt [7] show that all the standard types of games can be implemented in a manner that yields predictions consistent

with the Nash equilibrium for some parameter values, and yet in each case the observed data will shift dramatically in response to a payoff change that does not alter the Nash prediction. Similar anomalous results have been reported in many other experiments, e.g., matching pennies games, centipede games, two-stage games, market pricing games, and bargaining games [8-12]. We will present the main argument of a *social dilemma game* for which Nash's theory predicts a unique equilibrium that is "bad" for all concerned, and a *coordination game* in which any common effort level is an equilibrium, i.e. the Nash equilibrium makes no prediction at all.

The social dilemma is based on a story in which two travelers lose luggage with identical contents, and the airline official promises to pay any claim in an acceptable range as long as the claims are equal. If not, the person making the higher claim is assumed to have lied, and both will be reimbursed at the lower claim, with a reward, $R > 1$, being deducted from the reimbursement to the high claimant and given to the low claimant. A *Nash equilibrium* in this context is a pair of claims that survives an "announcement test:" if each person writes in their claim and then announces it as they turn in their claim sheet, neither should want to reconsider. Since the travelers file their claims separately, each will have a temptation to "undercut" any agreed on common claim. For example, suppose the range of acceptable claims is from 80 to 200, with a reward parameter, R , equal to 10. A common claim of 200 yields 200 for both, but a deviation by one person to 199 would profitably raise that person's payoff to $199 + 10$. The incentive to undercut the other's decision by 1 implies that the maximum claim of 200 is never an optimal choice, irrespective of the beliefs one has about the other's claim choice. Consequently, a rational person must assign zero probability to a choice of 200. But once 200 is ruled out as a possibility, 199 can be ruled out on the same grounds, and this logic can be repeated until the only beliefs rational players can have are that claims will be 80. In fact, 80 is the unique Nash equilibrium, despite the fact that both would be better off by claiming a high amount.

The paradoxical outcome of this "traveler's dilemma" game was first derived by Basu [3]. He did not expect behavior to converge to the Nash prediction for low values of R , but as he noted, none of the standard modifications of game theory can predict this anticipated deviation. Capra, Goeree, Gomez, and Holt [14] conducted an experiment based on this game form, using randomly matched student subjects who made claim decisions independently in a sequence of ten

periods. Earnings ranged from \$24 to \$44 and were paid in private, immediately after the experiment. With $R = 50$, the average claim was quite close the Nash prediction of 80 in the final 5 rounds, but with $R = 10$, the average claim started high (at about 180) and moved away from the Nash prediction, ending up at 186 for the last five rounds. The frequency of actual decisions for the final five rounds is indicated in figure 1 by the blue bars for $R = 50$ and by the red bars for $R = 10$. The yellow bars show the frequency of decisions for an intermediate treatment with $R = 25$. The task for theory is to explain these treatment differences, which sharply contrast the Nash prediction of 80, independent of R .




Figure 1. Claims in a Traveler's Dilemma with $R = 50$ (blue), $R = 25$ (yellow), and $R = 10$ (red).

The second game has a similar structure, with payoffs again being determined by the minimum of the two players' decisions. In this game, the decisions are "effort levels," and the joint production process is such that it requires both players to perform a costly task in order to

raise the level of production. The payoff for each player is the minimum of the two efforts, minus the cost of the player's own effort: $\pi_i = \min\{x_1, x_2\} - cx_i$, where x_i is player i 's effort level and $c < 1$ is a cost parameter. Just as in the traveler's dilemma, both players would be better off if they could coordinate on high decisions. The difference is that in a coordination game such a high-effort outcome *is* an equilibrium. Indeed, consider a common effort level, and notice that a unilateral increase is costly but does not affect the minimum. Similarly, a unilateral decrease from any common effort level will reduce the minimum by more than the cost saving, since $c < 1$. Hence, *any* common effort level would survive the Nash announcement test, so the maximum effort is an equilibrium, as is any other level.

While a low effort cost makes it relatively safe to choose a high effort, a high cost makes this action risky as it may not be matched by the other player, which suggests that actual behavior might be sensitive to changes in the cost parameter. Goeree and Holt [15] report an experiment with randomly matched pairs of subjects who made effort choices in the range from 110 to 170 (pennies). With a relatively low effort cost of $c = .25$, efforts tended to be in the upper part of this range, and with a high effort cost of $c = .75$, the data pattern was reversed. The time-sequences of average effort choices for three groups of 10 subjects in each treatment are shown in figure 2, with an upward pattern for the low-effort cost treatment (thin blue lines) and an essentially symmetric downward adjustment for the high-effort cost treatment (thin green lines). The thick blue and green lines show average efforts for each treatment (the red lines are discussed below). The strong treatment effect, which is consistent with simple intuition about the effects of effort costs, is not predicted by the Nash equilibrium or any standard variant of it. Another interesting result of the experiment is that subjects may "get stuck" in a low-effort equilibrium, even though there is there is a high-effort outcome that is not only better for all concerned, but it is also a Nash equilibrium [16-17].

To summarize, experimental data for both games reveal that the most salient feature of observed behavior is not predicted by classical game theory, where the assumption is that the decision with the highest expected payoff will invariably be chosen, no matter how small the payoff difference. In practice, players are usually unsure about what others will do, and there may be randomness in the way individuals accumulate, process, and act on strategic information.

In the coordination game with a low cost, for example, the potential losses from "overshooting" an optimal decision are much less than from "undershooting." Thus any noise in the decision making process might lead players to raise their effort decisions, which can further raise the payoffs associated with higher decisions. In an interactive, strategic situation, this process can "snowball" leading to a large increase in decisions. Even relatively small payoff asymmetries and small amounts of noise can have a large effect in an interactive, strategic game in which these effects are compounded endogenously, and this intuition is a key element of the static and dynamic theories presented in the next two sections.




Figure 2. A Coordination Game: Data Averages (blue and green lines) and Evolutionary Predictions (red lines).

Evolutionary Dynamics

A natural approach to explaining divergent adjustment patterns like those in figure 2 is to develop models of evolution and learning. The idea behind evolution in this context is not based on the biological notion of reproductive fitness, but rather that people tend to adopt strategies that have

worked well in the past. Besides incorporating payoff-seeking behavior, evolutionary models typically add in some stochastic elements that are reminiscent of mutation in biological models. For example, Anderson, Goeree, and Holt [18] present a model that has its roots in statistical physics and that can be seen as quantification of "learning direction theory" [19]. Individuals drawn from a population of players are characterized by their decision, $x(t)$, at time t , which earns an expected payoff of $\pi^e(x(t), t)$ where the expectation is taken with respect to the current population distribution of decisions, $F(x, t)$. The evolutionary assumption is that a player's decision tends to increase if the expected payoff function is increasing at $x(t)$, and vice versa. So the decisions evolve over time, with the rate of change proportional to the slope of expected payoff, π^e , plus a stochastic Brownian motion term: $dx = \pi^e(x(t), t)dt + \mu dw(t)$, with μ an error parameter that determines the relative importance of the random shock $dw(t)$. Techniques from classical physics can be used to show that this stochastic adjustment process translates into the following differential equation for the population distribution:

$$\frac{\partial F(x, t)}{\partial t} = -\pi^e(x, t) f(x, t) + \mu f'(x, t), \quad (1)$$

where f and f' represent the population density and its slope. This is the famous Fokker-Planck equation which is derived in this interactive context in [18]. The intuition for the two terms on the right side of (1) is that, when payoffs are increasing, probability mass will move the right (lowering F), and when the density is increasing, the effect of noise is to "flatten" the density pushing probability mass to the left (raising F).

The expected payoff function in (1) depends on the distribution of decisions, so the evolutionary process in (1) is in general a non-linear partial differential equation and much more complex than it first appears. Nevertheless, we can solve for the time paths of the distribution function numerically, given an initial distribution of efforts. The red lines in figure 2 show the trajectories of average efforts for the two treatments of the coordination game experiment. This model explains both the strong treatment effect and the general qualitative features of the adjustment path. Alternatives to the evolutionary model just described can be found in [20-27].

Learning Dynamics

Evolutionary models, when used to describe laboratory data, are sometimes criticized on the grounds that they ignore the cognitive abilities of human subjects. This critique is less severe than it seems at first, as a number of researchers have pointed out the close connections between standard replicator models of biological evolution and more cognitive learning models. The learning model that is closest to the evolutionary approach is "reinforcement learning" based on the psychological insight that successful strategies will be reinforced and used more frequently. Roth and Erev [12, 28] formalize these insights by giving each decision an initial weight, and then adding the payoff actually obtained for a decision chosen to the weight for that decision. The model is stochastic in the sense that the probability that a decision is taken is the ratio of its own weight to the sum of all weights. Simulations show that reinforcement learning models explain key features of data from many economics experiments [28].

Alternatively, learning can be modeled in terms of beliefs about others' decisions. For example, suppose that beliefs are characterized by a weight for each possible value of the other player's decision, and that the subjective probability associated with each decision is its weight divided by the sum of all weights. If the Bayesian prior distribution over decisions is Dirichlet and the generating process is thought to be multinomial, then the posterior is Dirichlet with the updated weights determined by adding 1 to the weight of the decision that was observed while the other weights are unaltered. The resulting subjective probabilities determine expected payoffs, which in turn determine players' choice probabilities. To allow for some randomness in decision making, a probabilistic choice rule can be used to inject some noise into the system. These rules have the property that a decision with a higher expected payoff is more likely to be chosen although not necessarily with probability one. For instance, the familiar "logit" form that is widely used in empirical work assumes that the probability of selecting a particular decision, i , is an exponential function of its expected payoff, $\pi^e(i)$:

$$Pr(i) = \frac{\exp(\pi^e(i)/\mu)}{\sum_{j=1}^N \exp(\pi^e(j)/\mu)}, \quad i = 1, \dots, N, \quad (2)$$

where the denominator ensures that the probabilities for the N decisions sum to one, and the error parameter μ determines how sensitive choice probabilities are to payoff differences. As μ goes to infinity, all choice probabilities are equal, regardless of payoff differences, but small payoff differences will have large effects when μ goes to zero.

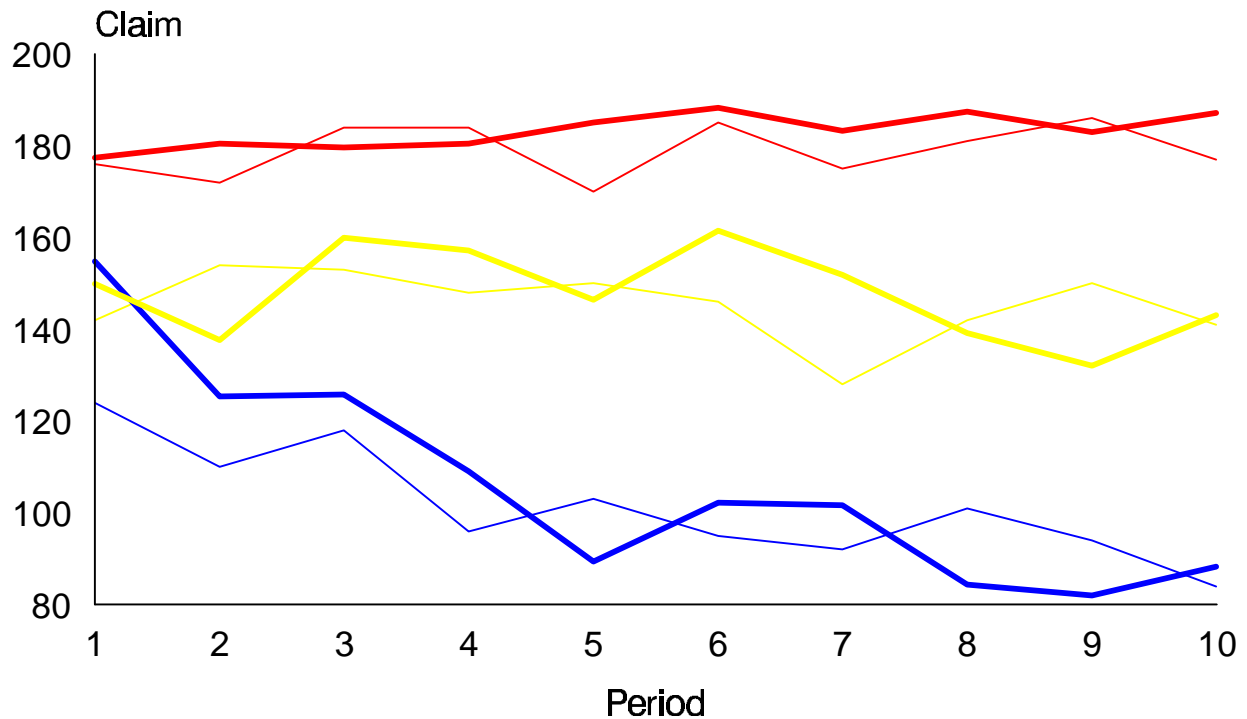


Figure 3. Adjustment Patterns in the Traveler's Dilemma: Human Data (thick lines) and Simulated Bayesian Learning (thin lines), by Color-Coded Treatments.

Computer simulations based on this noisy learning model begin with a specification of initial beliefs (often uniform), which determine the expected payoffs and choice probabilities in (2). Simulated players are matched and see the other player's randomly determined decision, which is then used to update the weights that determine beliefs, and hence expected payoffs. Then a new round of simulated decisions are generated with (2), and the process continues. Figure 3 shows computer simulations and actual human data for groups of ten randomly matched players in the three treatments of the traveler's dilemma game. Simulations of this type have also been

used to explain adjustment patterns in experiments based on signaling games [29], coordination games [15], and a price competition game [11]. At this time, there is no consensus about which dynamic model is best, and the work of Camerer and Ho [30] suggests that a hybrid formulation might provide the best fits to laboratory data. See [31] for a survey.

Logit Equilibrium

In a steady state, the distribution function in (1) stops evolving, i.e., $dF/dt = 0$. Setting the right side of (1) to zero, dropping the time arguments, and rearranging yields a differential equation in the equilibrium density: $f'(x)/f(x) = \pi^e(x)/\mu$. This equation can be integrated to yield $f(x) = k \exp(\pi^e(x)/\mu)$ which is the continuous analogue of (2). Thus a gradient-based evolutionary model with Brownian motion produces the logit model in a steady state. Likewise, when simulated behavior in learning models with logit errors settles down, the steady state distributions of decisions are consistent with the logit rule in (2).

In equilibrium, condition (2) must be interpreted carefully, since the choice density determines the expected payoffs, which in turn determine the choice density. In other words, the *logit equilibrium* is a fixed point: the "belief" density that goes into the expected payoff function on the right side of (2) must match the "choice" density that comes out of the logit rule on the left side. McKelvey and Palfrey [32-33] prove that such a fixed point exists for all games with a finite number of decisions. While the logit form is computationally convenient, it is possible to use other probabilistic choice rules [34-35], and McKelvey and Palfrey's proof covers the general case. Their elegant proof is based on a fixed-point theorem, like Nash's half-page Nobel Prize winning proof that was published in these *Proceedings*.

The equilibrium conditions in (2) can be solved numerically using an error parameter that is estimated from the data using standard maximum likelihood methods. Figure 4 shows the logit equilibrium densities for the treatments of the traveler's dilemma game. Note that the theoretical densities pick up the general location of the data frequencies in figure 1, with the colors used to match predictions and data for each treatment. There are certainly discrepancies, such as the underprediction of claims in the intermediate $R = 25$ treatment, but the logit equilibrium predictions are obviously superior to the Nash prediction of 80, which is independent of the

reward parameter R . The logit equilibrium has been successfully applied to explain behavior in a variety of environments including coordination games, signaling games, abstract matrix games, public goods games, jury decision making, and rent-seeking contests [32-33, 36-39].




Figure 4. Logit Equilibrium Densities for the Traveler's Dilemma with $R = 50$ (blue), $R = 25$ (yellow), and $R = 10$ (red).

A Model of Iterated Noisy Introspection

We have discussed learning and equilibrium in the previous two sections, but these techniques will not work when there is no opportunity for players to learn, adjust strategies, and reach an

equilibrium. "One-shot" games are particularly appropriate for modeling many political contests, legal disputes, special auctions, or other strategic situations that are not repeated.

The only way to try to avoid surprises in these situations is to use general intuition and introspection about what the other player(s) might do, what they think others might do, etc. Edgar Allen Poe mentions this type of iterated introspection in an account of a police inspector trying to decide where to search for the "purloined letter." Keynes [40] describes iterated introspection in the context of investors who want to purchase stocks that others will decide to purchase subsequently, because they think others will purchase, etc., causing the price to appreciate. Economists have long thought about such iterated expectations, and have noted that the resulting "infinite regress" often leads to a Nash equilibrium, and hence is of limited relevance in describing a new approach to behavior in games played only once.

Naturally, our approach is to introduce noise into the introspective process, with the intuition that players' own decisions involve less noise than their perceptions of others' decisions, which in turn involves less noise than other players' perceptions of others' decisions, etc. Our model is based on the probabilistic choice mapping in (2), which we will express compactly as $p = \phi_\mu(q)$, where q represents the vector of belief probabilities that determine expected payoffs, and p is the vector of probabilities that is determined by the probabilistic choice rule ϕ_μ with error parameter μ . We will assume that the noise parameter associated with each higher level of iterated introspection is $t > 1$ times the error parameter associated with the lower level. For instance, $p = \phi_\mu(\phi_{t\mu}(q))$ represents a player's noisy (μ) response to the other player's noisy ($t\mu$) response to beliefs q . The "telescope" parameter t determines how fast the error rate blows up with further iterations; the error rate for the n th iteration is given by $t^{n-1}\mu$. We are interested in the choice probabilities in the limit as the number of iterations goes to infinity:

$$p = \lim_{n \rightarrow \infty} \phi_\mu(\phi_{t\mu}(\cdots \phi_{t^{n-1}\mu}(q))). \quad (3)$$

In [41] we use continuity arguments to show that this limit is well defined when $t > 1$. Since ϕ_∞ maps the whole probability simplex to a single point, the process is independent of the initial belief vector q . Goeree and Holt [41] show that (3) provides a good explanation of (non-equilibrium) play in many types of one-shot games; see also [42-43] for alternative approaches.

The logit equilibrium arises as a limit case of this two parameter introspective model. Recall that a logit equilibrium is a fixed point of ϕ_μ , i.e. a vector p^* that satisfies $p^* = \phi_\mu(p^*)$, and note that for $t = 1$, a fixed point of ϕ_μ is also a fixed point of (3). So, if the introspective model converges for $t = 1$, the result is a logit equilibrium (although, in general, convergence is only ensured for $t > 1$). To summarize, the logit equilibrium generalizes Nash by relaxing the assumption of perfect decision making, and the introspective model generalizes logit by relaxing the assumption of perfect consistency between actions and beliefs.

Conclusion

Game theory is the closest thing to a unifying theory in social science, and it evokes some of the strongest antagonism as well. Critics argue that people are not perfectly rational, and that the experimental support for game theory is mixed. Daniel Kahneman, a noted Princeton psychologist, remarked in a plenary address: "When an economist says the evidence is mixed, that means the theory says one thing and the data say something else." For most economic theorists, the subtext on this would be that there must be something wrong with the experiments because the theory is logically correct. The problem with this normative defense is that what is optimal in a game like the traveler's dilemma depends on what the other players actually do, not on what some theory says they should do.

This paper describes three complementary modifications of classical game theory. The models of introspection, learning/evolution, and equilibrium contain the common stochastic elements that represent errors or unobserved preference shocks. These three approaches are like the "three friends" of classical Chinese gardening (pine, prunus, and bamboo), they fit together nicely, each with a different purpose. Models of iterated noisy introspection are used to explain beliefs and choices in games played only once, where surprises are to be expected, and beliefs are not likely to be consistent with choices. With repetition, beliefs and decisions can be revised via learning or evolution. Choice distributions will tend to stabilize when there are no more surprises in the aggregate, and the resulting steady state constitutes a noisy (quantal response) equilibrium.

These theoretical perspectives have allowed us to predict initial play, adjustment patterns, and final tendencies in a series of laboratory experiments. Data patterns that our colleagues would

previously characterize as "behavioral" (i.e., consistent with intuition but not with theory) are being picked up by these new stochastic game-theoretic models. There are discrepancies and surprises, but the overall pattern of results is surprisingly coherent, especially considering that we are using human subjects in interactive situations. In fact, the coauthor with a second degree in physics (Goeree) sometimes remarks that he is getting "that old physics feeling" when something unexpected happens in an economics experiment.

Laboratory experiments have been intimately connected with the development of game theory, starting with the reaction to Nash's seminal theorem that appeared in this journal. Two of the three recipients of the first Nobel Prize in Economics given to game theorists (Nash and Selten) conducted experiments. Patterns of actual human data provide the landmarks that are needed to avoid becoming lost in the jungle of possibilities once theorists move away from assumptions of perfect rationality. The resulting models have the empirical content that makes them relevant for playing games, not just for doing theory.

References

1. Nash, J. (1950) *Proceedings of the National Academy of Sciences, U.S.A*, 36, 48-49.
2. Nasar, S. (1998) *A Beautiful Mind*, New York: Simon and Schuster.
3. Roth, A. E. (1995) in the *Handbook of Experimental Economics*, eds. Kagel, J. H. and Roth, A. E. (Princeton University Press, Princeton: New Jersey), pp. 3-109.
4. von Neumann, J. & Morgenstern, O. (1944) *Theory of Games and Economic Behavior*, (Princeton University Press: Princeton: New Jersey).
5. Selten, R. (1975) *International Journal of Game Theory*, 4, 25-55.
6. Harsanyi, J. C. (1967-1968) *Management Science*, 14, 159-182, 320-334, 486-502.
7. Goeree, J. K. & Holt, C. A. (1999) "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," working paper, University of Virginia.
8. Ochs, J. (1994) *Games and Economic Behavior*, 10, 202-217.
9. McKelvey, R. D. & Palfrey, T. R. (1992) *Econometrica*, 60, 803-836.
10. Beard, T. R. & Beil, R. O. (1994) *Management Science*, 40(2), 252-262.

11. Capra, C. M., Goeree, J. K., Gomez, R. & Holt, C. A. (1999) "Learning and Noisy Equilibrium Behavior in an Experimental Study of Imperfect Price Competition," working paper, University of Virginia.
12. Roth, A. E. & Erev, I. (1995) *Games and Economic Behavior*, 8, 164-212.
13. Basu, K. (1994) *American Economic Review*, 84(2), 391-395.
14. Capra, C. M., Goeree, J. K., Gomez, R. & Holt, C. A. (1999) "Anomalous Behavior in a Traveler's Dilemma?" forthcoming in the *American Economic Review*.
15. Goeree, J. K. & Holt, C. A. (1998) "An Experimental Study of Costly Coordination," working paper, University of Virginia.
16. Van Huyck, J. B., Battalio, R. C. & Beil, R. O. (1990) *American Economic Review*, 80, 234-248.
17. Cooper, R., DeJong, D. V., Forsythe, R. & Ross, T. W. (1992) *Quarterly Journal of Economics*, 107, 739-771.
18. Anderson, S. P., Goeree, J. K. & Holt, C. A. (1999) "Stochastic Game Theory: Adjustment and Equilibrium with Bounded Rationality," working paper, University of Virginia.
19. Selten, R. & Buchta, J. (1994) "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions," University of Bonn Discussion Paper B-270.
20. Foster, D. & Young, P. (1990) *Theoretical Population Biology*, 38, 219-232.
21. Friedman, D. & Yellin, J. (1997) "Evolving Landscapes for Population Games," University of California, Santa Cruz, draft.
22. Fudenberg, D. & Harris, C. (1992) *Journal of Economic Theory*, 57(2), 420-441.
23. Kandori, M., George, M. & Rob, R. (1993) *Econometrica*, 61(1), 29-56.
24. Binmore, K., Samuelson, L. & Vaughan, R. (1995) *Games and Economic Behavior*, 11, 1-35.
25. Crawford, V. P. (1991) *Games and Economic Behavior*, 3, 25-59.
26. Crawford, V. P. (1995) *Econometrica*, 63, 103-144.
27. Young, P. (1993) *Econometrica*, 61, 57-84.
28. Erev, I. & Roth, A. E. (1998) *American Economic Review*, 88(4), 848-881.
29. Cooper, D. J., Garvin, S. & Kagel, J. H. (1994) *RAND Journal of Economics*, 106(3), 662-

683.

30. Camerer, C. & Ho, T. H. (1999) forthcoming in *Econometrica*.
31. Fudenberg, D. & Levine, D. K. (1998) *Learning in Games* (MIT Press, Cambridge, MA).
32. McKelvey, R. D. & Palfrey, T. R. (1995) *Games and Economic Behavior*, 10, 6-38.
33. McKelvey, R. D. & Palfrey, T. R. (1998) *Experimental Economics*, 1(1), 1-41.
34. Rosenthal, R. W. (1989) *International Journal of Game Theory*, 18, 273-292.
35. Chen, H. C., Friedman, J. W. & Thisse, J. F. (1997) *Games and Economic Behavior*, 18, 32-54.
36. McKelvey, R. D. & Palfrey, T. R. (1998) "An Experimental Study of Jury Decisions," working Paper, Caltech.
37. McKelvey, R. D., Palfrey, T. R. & Weber, R. A. (1999) "The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2×2 Games with Unique Mixed Strategy Equilibria," working Paper, Caltech.
38. Anderson, S. P., Goeree, J. K. & Holt, C. A. (1998) *Journal of Political Economy*, 106(4), 828-853.
39. Anderson, S. P., Goeree, J. K. & Holt, C. A. (1998) *Journal of Public Economics*, 70, 297-323.
40. Keynes, J. M. (1965) *The General Theory of Employment, Interest, and Money* (Harcourt, Brace, and World, Inc., New York).
41. Goeree, J. K. & Holt, C. A. (1999) "Models of Noisy Introspection," working paper, University of Virginia.
42. Capra, C. M. (1998) "Noisy Expectation Formation in One-Shot Games: An Application to the Entry Game," working paper, University of Virginia.
43. Harsanyi, J. C. & Selten, R. (1988) *A General Theory of Equilibrium Selection in Games* (MIT Press, Cambridge, MA).