

Theory and experiment in the analysis of strategic interaction

Vincent P. Crawford

One cannot, without empirical evidence, deduce what understandings can be perceived in a nonzero-sum game of maneuver any more than one can prove, by purely formal deduction, that a particular joke is bound to be funny.

Thomas Schelling, *The Strategy of Conflict*

1 INTRODUCTION

Much of economics has to do with the coordination of independent decisions, and such questions – with some well-known exceptions – are inherently game theoretic. Yet when the Econometric Society held its First World Congress in 1965, economic theory was still almost entirely non-strategic and game theory remained largely a branch of mathematics, whose applications in economics were the work of a few pioneers. As recently as the early 1970s, the profession's view of game-theoretic modeling was typified by Paul Samuelson's customarily vivid phrase, "the swamp of n -person game theory"; and even students to whom the swamp seemed a fascinating place thought carefully before descending from the high ground of perfect competition and monopoly.

The game-theoretic revolution that ensued altered the landscape in ways that would have been difficult to imagine in 1965, adding so much to our understanding that many questions whose strategic aspects once made them seem intractable are now considered fit for textbook treatment. This process was driven by a fruitful dialogue between game theory and economics, in which game theory supplied a rich language for describing strategic interactions and a set of tools for predicting their outcomes, and economics contributed questions and intuitions about strategic behavior

against which game theory's methods could be tested and honed. As game-theoretic formulations and analyses enriched economics, economic applications inspired extensions and refinements of game theory's methods, transforming game theory from a branch of mathematics with a primarily normative focus into a powerful tool for positive analysis.

To date this dialogue has consisted mostly of conversations among theorists, with introspection and casual empiricism the main sources of information about behavior. A typical exchange proceeds by modeling an economic environment as a non-cooperative game; identifying its equilibria; selecting among them as necessary using common sense, equilibrium refinements, dynamic arguments, or convenience; comparing the selected equilibrium with stylized facts and intuitions about outcomes; and eliminating discrepancies, as far as possible, by adjusting the model or proposing new selection criteria. The unstated goal of most such analyses has been to predict behavior entirely by theory.

Although this approach has plainly been productive, it has also revealed the limits of what can be learned by theory alone. Theoretical analyses (traditional or adaptive) usually yield definite predictions only under strong assumptions, which are reasonable for some applications but unrealistic and potentially misleading for many others. As a result, most strategic applications raise questions about the principles that govern behavior that are not convincingly resolved by theory, in addition to questions about preferences and the environment like those encountered in non-strategic applications. Further progress in understanding those principles now seems likely to depend as much on systematic observation and careful empirical work as on further advances in theory.

Experiments will play a leading role in this empirical work. Behavior in games is notoriously sensitive to the details of the environment, so that strategic models carry a heavy informational burden, which is often compounded in the field by an inability to observe all relevant variables. Important advances in experimental technique over the past three decades allow a control that often gives experiments a decisive advantage in identifying the relationship between behavior and the environment.¹ There is now a substantial body of experimental work that uses well-motivated subjects and careful designs to address central questions about strategic behavior. I believe this work deserves to be taken seriously. For many questions it is the most important source of empirical information we have, and it is unlikely to be less reliable than casual empiricism or introspection. More generally, I believe that there is much to be gained by supplementing conversations among theorists with a dialogue between theorists and experimentalists, in which theoretical ideas are confronted with observation as well as intuition.

This chapter considers the roles of theory and experiment in the analysis of strategic interaction, with the goal of encouraging and focusing the dialogue that has already begun.² I emphasize the benefits to theorists of thinking about experiments, which is both what I know best and the direction in which the dialogue seems most in need of encouragement. My principal goals are to identify the kinds of theory that are useful in interpreting experimental evidence and to draw out the conclusions about behavior the evidence suggests. Accordingly, the discussion is organized along strategic rather than economic lines, even though this cuts across conventional boundaries in the experimental literature; and I favor experiments that seek clear identification of general principles, even when this comes at the expense of realism.³ This approach makes applications seem more remote, but it exploits the generality of game-theoretic formulations in a way that seems most likely to yield the depth of understanding the analysis of economic models requires.

The experimental evidence suggest that none of the leading theoretical frameworks for analyzing games – traditional non-cooperative game theory, cooperative game theory, evolutionary game theory, and adaptive learning models – gives a fully reliable account of behavior by itself, but that most behavior can be understood in terms of a synthesis of ideas from those frameworks, combined with empirical knowledge in proportions that depend in predictable ways on the environment. In this view theory and experiment have complementary roles, with theory providing a framework within which to gather and interpret the empirical information needed to close the model, in addition to developing its implications, and experiments mapping the boundaries of the environments and aspects of behavior for which theoretical ideas allow adequate predictions, and identifying and observing the aspects of behavior theory does not reliably determine.

The chapter is organized as follows. Section 2 reviews the leading theoretical frameworks and unresolved questions. Section 3 gives an overview of experimental designs. Sections 4–6 discuss experimental evidence, and section 7 is the conclusion.

2 THEORETICAL FRAMEWORKS AND UNRESOLVED QUESTIONS

The leading theoretical frameworks for analyzing behavior in games – traditional non-cooperative and cooperative game theory, evolutionary game theory, and adaptive learning models – reflect different views of how beliefs and/or strategies are determined, each of which has something to contribute to our understanding of experimental results. This section

reviews them, emphasizing important aspects that may be unfamiliar and concluding with a discussion of unresolved questions.

In traditional game theory behavior in a game is determined entirely by its *structure*, which consists of its players, the decisions they face and the information they have when making them, how their decisions determine the outcome, and their preferences over outcomes. The structure incorporates any repetition, correlating devices, or opportunities for communication. Some theories allow behavior to be influenced by other factors, such as how the game is presented or the social setting; I call such factors the *context*.

A player's decisions are summarized by a complete contingent plan called a *strategy*, which specifies his decision as a function of his information at each point at which he might need to make one. Players' strategies should be thought of as chosen simultaneously, at the start of play; taken together they determine an outcome in the game.

Something is *mutual knowledge* if all players know it, and *common knowledge* if all players know it, all players know that all players know it, and so on ad infinitum.

The essential difficulty of game theory is that the consequences of players' decisions depend on decisions by others that they cannot observe, and must therefore predict. In all but the simplest games, players typically bear uncertainty about each other's strategies, which I shall call *strategic uncertainty*. To focus on the issues strategic uncertainty raises, I simplify the problem of characterizing individual decisions by adopting the standard assumption that it is mutual knowledge that players are *rational* in the sense that their expectations about each other's strategies can be summarized by probability distributions called *beliefs*, and their preferences over uncertain outcomes can be described by assigning numerical *payoffs* to outcomes so that they maximize expected payoffs, given their beliefs.⁴

Strategic sophistication refers to the extent to which a player's beliefs and behavior reflect his analysis of the environment as a game rather than a decision problem, taking other players' incentives and the structure into account.⁵ Like strategic uncertainty it is a multi-dimensional concept, which must be adapted to specific settings as illustrated below.

2.1 Traditional non-cooperative game theory

Traditional non-cooperative game theory is distinguished by the use of Nash's notion of equilibrium to describe players' behavior throughout the analysis. An *equilibrium* is a combination of strategies such that each player's strategy maximizes his expected payoff, given the others'. It reflects

self-confirming beliefs in that rational players will choose equilibrium strategies if – and in general only if – they correctly anticipate each other's choices. This result can be formalized as follows, taking a broader, beliefs-based interpretation of equilibrium that is useful. Assume that rationality and the structure are mutual knowledge; that players have a common prior, so that any differences in their beliefs can be traced to differences in information; and that their beliefs are common knowledge. Then any two players' beliefs about a third player's strategy must be the same and these common beliefs, viewed as mixed strategies, must be in equilibrium (Aumann and Brandenburger (1995)).⁶ In this *equilibrium in beliefs*, a player's mixed strategy represents other players' beliefs about his realized pure strategy, about which he himself need not be uncertain, and players' beliefs determine their optimal strategies and expected payoffs. (Assuming that each player bears the same uncertainty about his realized pure strategy as other players yields the standard notion of equilibrium in strategies.)

The stated conditions are the weakest sufficient conditions available for games in general. Equilibrium normally requires, in addition to rationality, the assumption that players' beliefs are coordinated on the same outcome. In applications this is either assumed, with beliefs taken as given, or viewed as the result of independent predictions based on a common *coordinating principle*, such as a convention, norm, or focal point; an equilibrium refinement; or a complete theory of equilibrium selection (Harsanyi and Selten (1988)).⁷ Thus, traditional equilibrium analysis assumes an extreme form of strategic sophistication, in that players must understand the structure and how their partners will respond well enough to make beliefs or strategies mutual knowledge, eliminating strategic uncertainty. This assumption is appropriate for settings simple or familiar enough that players can predict each other's responses, and it is often helpful in thinking about players' likely responses to entirely new environments. However, it is plainly too strong for many applications. Yet assuming only common knowledge of rationality and the structure, with no restrictions on beliefs, implies only the iterated elimination of strategies that are never weak best replies, which in many games yields no useful restrictions on behavior. To analyze such games one must impose restrictions on beliefs or behavior from other sources.

I call a coordinating principle *structural* if it depends entirely on the structure of the game, and *contextual* if it also depends on the context. A principle is *inductive* if it predicts behavior directly from behavior in analogous games, and *deductive* if it is defined on a more general class of games, and predicts behavior in the current game only indirectly.⁸ Traditional game theory usually studies principles that are structural and

deductive. However, this is a matter of custom rather than logic, and beliefs can be coordinated equally well by contextual or inductive principles. Such principles often play important roles in experiments because they place more realistic demands on subjects' information and subjects find direct analogies more convincing than abstract arguments.

2.2 Cooperative game theory

Cooperative game theory studies frictionless bargaining among rational players who can make binding agreements about how to play a game. Like non-cooperative game theory, it is structural and assumes an extreme form of strategic sophistication. It differs in three ways: (i) it summarizes the structure by the payoffs players can obtain acting alone or in coalitions, suppressing other aspects; (ii) instead of explicitly modeling players' decisions, it assumes that they reach an efficient agreement; and (iii) it uses simple symmetry or coalition rationality assumptions to characterize how players share the resulting surplus. These features give cooperative game theory a strong comparative advantage in analyzing behavior in environments whose structures cannot be observed or described precisely.

2.3 Evolutionary game theory

Evolutionary game theory studies environments in which games are played repeatedly in populations, analyzing the dynamics of the population strategy frequencies under simple assumptions about how they respond to current expected payoffs. Although evolution presumably has little direct influence on behavior in experiments, evolutionary models are good templates for models of learning dynamics because they have interaction patterns like most experimental designs, they provide a framework for analyzing the effects of how players' roles and strategies are distinguished, and they suggest useful characterizations of the effects of strategic uncertainty. An evolutionary analysis is usually the first step toward understanding the dynamics of subjects' behavior, and combining the appropriate "evolutionary" structure with a realistic characterization of individual learning often yields a model well suited to describing experimental results.

In the simplest evolutionary models, a large population of players repeatedly play a symmetric game. I call the game that is repeated the *stage game* and strategies in the stage game *actions*, reserving "game" and "strategy" for the repeated game. Players are identical but for their actions. Their roles in the stage game are not distinguished, but their actions have a fixed common labeling, which gives meaning to statements like "players i

and j played the same action” or “player i played the same action in periods s and t .” Individual players play only pure actions, with payoffs determined by their own actions and the population action frequencies. This specification allows many of the symmetric interaction patterns studied in economics, including the familiar case of random pairing to play a two-person game (in which the stage game describes the simultaneous interaction of the entire population, with payoffs evaluated before the uncertainty of pairing is resolved).

In biology the law of motion of the population action frequencies is derived, usually with a functional form known as the replicator dynamics, from the assumption that players inherit their actions unchanged from their parents, whose reproduction rates, or *fitnesses*, are proportional to their payoffs (Maynard Smith (1982)). In economics similar dynamics are derived from plausible assumptions about individual adjustment (Schelling (1978, pp. 213–43), Crawford (1989, 1991), Friedman (1991)). The usual goal is to identify the locally stable steady states of the dynamics. A remarkable conclusion emerges: if the dynamics converge, they converge to a steady state in which the actions that persist are optimal in the stage game, given the limiting action frequencies; thus, the limiting frequencies are in Nash equilibrium.⁹ Even though players’ actions are not rationally chosen – indeed, not even chosen – the population collectively “learns” the equilibrium as its frequencies evolve, with selection doing the work of rationality and strategic sophistication.

In the Intersection and Confrontation examples of Crawford (1991, section 3), a large population of identical players are randomly paired to play games with common action labelings but undistinguished roles. In Intersection two drivers meet on different roads at an intersection and choose simultaneously between actions labeled Go and Stop, with payoffs of 1 if they choose different actions and 0 if they choose the same actions. Evolutionary dynamics converge to a frequency of Go of $1/2$ for any initial frequencies between 0 and 1, because Stop’s expected payoff exceeds Go’s if and only if the frequency of Go exceeds $1/2$. This outcome corresponds to the inefficient symmetric mixed-strategy equilibrium. In Confrontation two drivers confront each other on the same road and choose between actions labeled Left and Right, with payoffs of 1 if they choose the same actions and 0 if they choose different actions. The dynamics then converge to one of the frequencies of Right, 0 or 1, that corresponds to an efficient pure-strategy equilibrium; and the frequency $1/2$ that corresponds to the symmetric mixed-strategy equilibrium is unstable. In this case then the dynamics exhibit a simple form of history dependence in that the limiting equilibrium is determined by the initial frequencies. This and the more complex forms of history dependence in related models of learning dynamics play important

roles in describing the results of some experiments, illustrated in sections 6.1 and 6.3.

An evolutionary analysis can yield different outcomes in these games, even though their structures are identical, because in Intersection, but not in Confrontation, efficient coordination requires that undistinguished players choose actions with different labels. Similar differences in labeling often have substantive consequences in experiments because the labels are the language in which subjects interpret their experience, and in which inductive coordinating principles must sometimes be expressed. As the examples illustrate, evolutionary game theory has a system for modeling the effects of such differences. In Intersection the frequencies of the two efficient pure-strategy equilibria cannot even be represented in the state space used to analyze the dynamics, because the theory models the impossibility of systematic differences in aggregate action frequencies across roles that players cannot distinguish by assuming that undistinguished roles are filled by independent random draws from the same population.¹⁰ This device is easily extended to adaptive learning models with “evolutionary” structures, where it suggests a characterization of the effects of strategic uncertainty whose usefulness is illustrated in section 6.3.

Most discussions of evolutionary games in economics treat them as synonymous with random pairing, but many important applications are better modeled by assuming that the entire population plays a single n -person game. The same methods can be used to analyze the population dynamics in such games, known in biology as *games against the field*.¹¹ In the simplest such environments, a population of identical players repeatedly plays a symmetric stage game with undistinguished roles, one-dimensional action spaces, and common action labels. Each player's payoffs are determined by his own action and a summary statistic of all players' actions, such as the mean, minimum, or median.

In the Stag Hunt example of Crawford (1991, section 3), n players simultaneously choose between two efforts, 1 and 2. Their efforts yield a total output of $2n$ times the minimum effort, which they share equally; and the unit cost of effort is 1. Thus if all players choose the same effort their output shares more than repay the cost, but if anyone shirks the balance of the others' efforts is wasted. For any n , Stag Hunt has two symmetric pure-strategy equilibria, one in which all choose 2 and one in which all choose 1. Both of these equilibria are steady states. The same conclusions hold for the game in which players are randomly paired from a population of n to play two-person versions of Stag Hunt. Crawford (1991, figure 1) graphs the expected payoffs of efforts 1 and 2 against the population frequency of effort 1 for Stag Hunt with random pairing and against the field. With random pairing both equilibria are evolutionarily stable, and the

sets of initial frequencies from which the population converges to them – their *basins of attraction* – are equally large. Against the field, only the “all-1” equilibrium is stable, and its basin of attraction is almost the entire state space; other order statistics make the all-2 equilibrium locally stable, but with a small basin of attraction for order statistics near the minimum.

2.4 Adaptive learning models

Adaptive learning models describe players’ beliefs or strategies as the product of learning from experience with analogous games (Crawford (1989, 1991), Fudenberg and Kreps (1993), Marimon (1996)). The learning process is usually modeled as a repeated game, in which the analogies are transparent. The stage game is played either by a small all-inclusive group or in one or more populations, with “evolutionary” interaction patterns. Players’ actions and/or roles are distinguished by labels as in evolutionary game theory.

Adaptive learning is “adaptive” in that it need not be consistent with equilibrium in the stage game or the repeated game that describes the entire learning process.¹² Thus it allows for strategic uncertainty, often in arbitrary amounts. Players view actions as the objects of choice, and the dynamics of their choices are described either directly, or indirectly in terms of their beliefs, with actions modeled as best replies.¹³ Strategic sophistication is limited, with restrictions on behavior derived from simple, plausible assumptions about players’ adjustments or how they model each other’s behavior. These range from probabilistic responses to realized payoffs as in the psychological learning literature, which require no strategic sophistication at all (Andreoni and Miller (1995), Roth and Erev (1995)), to models like best-reply dynamics, fictitious play, and more general inertial dynamics, which require that players understand the structure but not other players’ decisions (Crawford (1995a), Broseta (1993a), Fudenberg and Kreps (1993)), and, finally, to models in which players have detailed models of other players’ decision processes, whose sophistication approaches that assumed in traditional analyses (Stahl and Wilson (1995)).

2.5 Unresolved questions

Well-informed experimental subjects usually exhibit some strategic sophistication, but often not enough to eliminate all strategic uncertainty before they begin to interact. Their beliefs are influenced by various kinds of coordinating principles, often contextual and inductive rather than structural and deductive. When beliefs are not perfectly coordinated at the start, learning typically yields rapid convergence to an equilibrium, in beliefs if

not in actions. However, the learning process is frequently history dependent, and strategic uncertainty, strategic sophistication, and the structure of learning rules often exert persistent influences on the outcome. Evolutionary and adaptive learning models, for instance, usually assume no strategic sophistication, but their dynamics do not always eliminate weakly dominated actions (Samuelson (1993)). Thus their predictions may be permanently off if players are sophisticated enough to eliminate such actions at the start.

The extent of strategic sophistication and strategic uncertainty, the coordinating principles that influence subjects' beliefs, and the structure of learning rules all appear to vary with the environment in predictable ways. There is a large body of experimental evidence on these patterns of variation from ultimatum and alternating-offers bargaining games and other dominance-solvable games, in which strategic sophistication is naturally identified with how many rounds of iterated deletion of dominated strategies players' beliefs reflect. There is also a large body of evidence from coordination and simultaneous-offers bargaining games and other games with multiple equilibria that survive iterated deletion of dominated strategies, where equilibrium requires what I shall call *simultaneous* coordination of beliefs and strategic sophistication can take more subtle forms. Sections 4 and 5 discuss evidence from these two kinds of environment that is "static," in that it can be understood without considering how behavior varies with repeated play. Section 6 considers "dynamic" evidence of both kinds.¹⁴

3 EXPERIMENTAL DESIGNS

This section discusses the designs used in most game experiments in economics. Many studies consider groups of related environments because the variations in behavior across them are often informative (Roth (1995a)). I call each such environment a *treatment*, a session in a treatment a *run*, and a set of related treatments an *experiment* (Roth (1994)).

A successful design must control the environment so that the results can be interpreted as responses to a clearly identified game. A typical design has one or more subject populations repeatedly playing a stage game in an "evolutionary" pattern (section 2.3), with the goal of testing theories of behavior in the stage game. Accordingly, the effects of repeated interaction are minimized by having subjects interact in small groups drawn from "large" populations, with repeated encounters unlikely or impossible; or in "large" groups with small influences on each other's payoffs.¹⁵ Subjects are usually told the outcome after each play, including their current partners' or all subjects' actions. To maintain control, communication and correlation

are allowed only as the stage game permits them. The stage game is otherwise free to vary, and can even be a repeated game. This freedom allows a wide range of strategic questions to be posed in tractable ways.

Subjects' unfamiliarity with such environments is overcome by using simple stage games and interaction patterns; explaining them in written instructions and question and answer sessions; and providing enough experience via practice rounds or repeated play to assure meaningful responses and reveal the effects, if any, of learning.

Non-strategic uncertainty is usually kept to a minimum to focus on strategic issues. Control of information is achieved by publicly announcing the structure at the start. The resulting condition, called *public knowledge*, comes as close as possible to inducing common knowledge in the laboratory.

Control over preferences is achieved by paying subjects according to their payoffs. Non-pecuniary effects are usually suppressed by avoiding frames with psychological associations and face-to-face or non-anonymous interactions (Roth (1995a, pp. 79–86)).¹⁶ Subjects' payments are normally linear functions of their game payoffs, with the results analyzed assuming risk neutrality. Sometimes, as in the "binary lottery" procedure of Roth and Malouf (1979), each subject is rewarded with a probability, again a linear function of his payoff, of winning a given amount of money (or the larger of two possible amounts). Under standard assumptions subjects then maximize the probability of winning the prize (or the larger prize), hence are risk neutral in a variable under experimental control.

Departures from these "consensus" designs are noted below only when they are important. Otherwise the designs can be assumed to involve one or more subject populations repeatedly playing a given stage game in an "evolutionary" pattern, with subjects motivated by one of the above methods, and with public knowledge.

4 DOMINANCE AND ITERATED DOMINANCE

This section discusses static evidence on dominance, iterated dominance, and closely related extensive-form refinements such as backward and forward induction. I begin with environments subjects seem to code as "abstract" rather than identifying them with games they are familiar with. I conclude with ultimatum and alternating-offers bargaining games.

4.1 Abstract games

Experiments with abstract games are well suited to studying strategic sophistication because they limit the effects of prior experience. Most work

in this area uses variants of two-person games like Stag Hunt or Battle of the Sexes, sometimes with outside options, in normal and/or extensive form. The conclusions are easy to summarize. Subjects avoid weakly or strongly dominated strategies, with frequencies usually greater than 90 percent. However, they rule out the possibility that others play dominated strategies with much lower frequencies, ranging from 20 percent to just over 80 percent; still fewer subjects rely on more than one round of iterated dominance; and the presence of dominated strategies often affects equilibrium selection even though they are rarely played (Beard and Beil (1994), Brandts and Holt (1993b), Cooper *et al.* (1994), Nagel (1995), Stahl and Wilson (1995), Van Huyck, Battalio, and Beil (1990, 1993)). Overall, subjects display significantly more strategic sophistication than evolutionary and adaptive learning models assume, but much less than is needed to justify many applications of iterated dominance and related refinements in economics.¹⁷

Beard and Beil (1994) investigated these phenomena more deeply by studying how outcomes vary with payoffs in two-person extensive-form games in which one player has a dominated strategy. They found that subjects' reliance on dominance varies in coherent, plausible ways with changes in the benefits to subjects and their partners, and in the cost a subject imposes on his partner by following the resulting strategy. They also found that experience in different roles made subjects more likely to rely on dominance in predicting the behavior of others in those roles.

In the elegant design of Nagel (1995), subjects simultaneously "guessed" numbers from 0 to 100, with the guess closest to a given fraction, p , of the population mean winning a prize. When $0 < p < 1$ this game has a unique equilibrium, which can be computed by iterated dominance: guesses greater than $100p$ are dominated; when these are eliminated guesses greater than $100p^2$ are dominated; and so on until (in the limit) only 0 remains. On the assumption that subjects ascribe a uniform level of sophistication to others, their initial responses reveal their levels of sophistication: a subject who thinks others guess randomly will guess $50p$; one who thinks others avoid dominated strategies but otherwise guess randomly will guess $50p^2$; and so on. Subjects never played equilibrium strategies; and most made guesses associated with only 1–3 rounds of dominance (see also Stahl (1994)).

Camerer *et al.* (1993) studied subjects' cognitive processes in a three-period alternating-offers bargaining game with a unique subgame-perfect equilibrium (assuming purely pecuniary payoffs), which is easily computed by backward induction (section 4.2). They used an ingenious computer interface that conceals the total payoffs of agreements in the three periods but allows subjects to look them up costlessly and as often as desired, but

only one at a time, while the computer automatically records their look-up patterns.¹⁸ If different cognitive processes yield different look-up patterns, the observed patterns allow direct tests of theories of cognition, along with their behavioral implications. This is an exciting prospect, which should speed progress in understanding strategic behavior.

Camerer *et al.* (1993) argued that backward induction in their game has a characteristic pattern in which: (i) subjects first check the third-period payoff, then the second-period payoff (possibly re-checking the third-period payoff), and finally the first-period payoff; (ii) most transitions are from later to earlier periods; and (iii) the most time is spent checking the second-period payoff.¹⁹ Aware that this is a larger (or at least different) leap of faith than most of us are used to, they remarked, “The reader may object to our characterization of the information search process that is inherent in equilibrium analysis. We are eager to hear alternative characterizations.” They also showed that a separate group of subjects, trained in backward induction and rewarded only for correctly computing their subgame-perfect equilibrium offers, came to exhibit just such a pattern.

As in related studies (section 4.2), subjects’ behavior was far from subgame-perfect equilibrium. Unlike with backward induction, subjects spent 60–75 percent of their time checking the first-period payoff, 20–30 percent checking the second-period payoff, and only 5–10 percent checking the third-period payoff, with most transitions from earlier to later periods. As expected, subjects who looked more often at the second- and third-period payoffs tended to make, or accept, initial offers closer to the subgame-perfect equilibrium; but there were no other clear correlations between look-up patterns and behavior. Despite Camerer *et al.*’s (1993) success in teaching subjects backward induction, repetition did not alter these patterns. Subjects’ focus on the first-period payoff, which determines the set of efficient agreements, suggests a concern for “fairness” of which we will see further evidence below.

4.2 Ultimatum and alternating-offers bargaining

The experimental literature on ultimatum and alternating-offers bargaining games with complete information is perhaps the largest body of evidence on dominance and iterated dominance (Roth (1995b), Camerer and Thaler (1995)). In these games two players, 1 and 2, take turns making offers about how to share a given “pie,” with player 1 going first. In the ultimatum game this process stops after player 1’s first offer, which player 2 must either accept or reject. Acceptance yields a binding agreement and rejection yields disagreement. In the alternating-offers game the process continues, indefinitely in some variants, until an offer is accepted, which

again yields a binding agreement. Rejection forces a delay of one period, which is costly because future agreements yield lower payoffs.

With purely pecuniary payoffs, the ultimatum game has a unique subgame-perfect equilibrium, in which player 1's first offer gives player 2 zero and player 2 accepts, yielding an efficient outcome.²⁰ The alternating-offers game also has a unique subgame-perfect equilibrium, in which player 1's first offer extracts all of player 2's surplus from accepting, given that player 2's best alternative is to make a counter offer one period later, chosen in the same way. In that equilibrium player 2 accepts, again yielding an efficient outcome.

The experimental results for both games are very different from these predictions. In ultimatum games first offers average 40 percent of the pie. In both games offers are rejected, with frequencies of 14–19 percent, and the frequency of inefficient delays and disagreements averages more than 25% (Forsythe, Kennan, and Sopher (1991, fn. 7, p. 261), Roth (1995b, table 4.5a, p. 293)). In alternating-offers games rejections are followed by “disadvantageous” counter offers that yield less than the rejected offer (and therefore violate dominance when payoffs are purely pecuniary), with frequencies of 65–88 percent (Roth (1995b, table 4.1, p. 265)).

Of particular interest are the parallel ultimatum experiments conducted in four countries by Roth *et al.* (1991). The results resemble the findings on offers and rejections just summarized, but with player 1s making systematically lower offers in two of the four countries. If the deviations from subgame-perfect equilibrium were due to lack of strategic sophistication, there would be no reason to expect the conditional rejection rates of player 2s to differ systematically across countries, so countries with lower offers should have more disagreements. Roth *et al.* (1991) found, instead, that rejection rates varied across countries in tandem with offers, so that countries with lower offers did not have more disagreements. In each country the modal offer in the tenth and final period maximized the expected payoffs of player 1s when their beliefs were estimated from that country's rejection rates.²¹

The frequency of rejections and disadvantageous counteroffers in ultimatum and alternating-offers experiments is often taken as evidence that subgame-perfect equilibrium requires too much sophistication to be descriptive, or that subjects' desire to be fair outweighs all strategic considerations. It is clear that subjects do not perceive their payoffs as purely pecuniary, even when these games are framed as abstractly as possible. Although there is some evidence that the required backward induction is too complex to describe behavior in alternating-offers games of more than two periods, the evidence from abstract games (section 4.1) suggests that behavior in ultimatum games is unlikely to be completely

unsophisticated. The simplest explanation of the results for ultimatum games one might hope for, then, is one in which player 1s are rational, motivated entirely by pecuniary payoffs, and respond in a strategically sophisticated way to the risk of rejection; and player 2s are rational but trade off pecuniary payoffs against their privately observed costs of accepting “unfair” offers, at a rate that may vary across countries, contexts, and players.

Adding this one plausible “epicycle” to the traditional model yields a parsimonious explanation of much of the evidence from ultimatum games. The behavior of player 1s is approximately consistent with equilibrium in beliefs, when beliefs are estimated from observed rejection rates. And the extended traditional model has the potential to explain other findings in which framing an ultimatum game so that player 1s “earned” the right to their roles, or allowing player 1s to impose an outcome without player 2s consent, moved outcomes closer to the subgame-perfect equilibrium (Roth (1995b)).

In the extended model players’ ideas about fairness are treated as exogenous non-pecuniary payoff parameters, whose distributions must be estimated empirically for each new environment, but which appear to vary across environments in stable, predictable ways. The resulting theory is a hybrid of traditional equilibrium analysis and standard econometric methods. Similar hybrids are important in environments discussed below.

5 SIMULTANEOUS COORDINATION

This section consider static evidence from games in which players make some decisions in ignorance of other players’ decisions and unaided by dominance. In such games equilibrium requires simultaneous coordination of beliefs, which relies on more detailed mental models of others’ decisions and more subtle forms of strategic sophistication.²² I begin with equilibrium refinements in signaling games. I then consider refinements, norms, and focal points in coordination games and unstructured bargaining games.

5.1 Signaling games

There is a small amount of static evidence on refinements in signaling games. Banks, Camerer, and Porter (1994) used the fact that the leading refinements – sequential equilibrium, the intuitive criterion, divinity, universal divinity, the never-a-weak-best-response criterion, and strategic stability – are nested, in the listed order, to construct a design that allows detailed comparisons of their performance in several games. The results

were consistent with some sequential equilibrium for 44–74 percent of the subject pairs in early periods and 46–100 percent in later periods. Each refinement predicted better than its coarser predecessor, up to and including divinity, but with success rates of at most 60 percent.

5.2 Coordination games

The only refinements that discriminate among the multiple strict equilibria in coordination games that have been tested experimentally are Harsanyi and Selten's (1988) notions of risk- and payoff-dominance and the "general theory of equilibrium selection" of which they are a part. Harsanyi and Selten's theory is of particular interest because, although they assume that players' beliefs and strategies converge to an equilibrium before play begins, the mental *tâtonnements* by which they model players' thought processes (the "tracing procedure" that underlies risk-dominance) are responsive to strategic uncertainty.

Perhaps the most informative tests of these notions to date are the experiments of Van Huyck, Battalio, and Beil (1990, 1991). They studied symmetric coordination games with structures like Stag Hunt, in which players without identified roles choose among seven "efforts," with payoffs determined by their own efforts and order statistics of all players' efforts. Here I focus on five leading treatments: one in which a game like Stag Hunt was played against the field by 14–16 subjects, with the order statistic the population minimum effort; one in which such games were played by 14–16 randomly paired subjects, with new partners each period and the order statistic the current pair's minimum effort; and three in which such a game was played against the field by nine subjects, with the order statistic the population median effort. In each case a player's payoff is highest, other things equal, when his effort equals the order statistic, so that any symmetric combination of efforts is an equilibrium. The equilibria are Pareto-ranked, with all preferring those with higher efforts; the highest-effort equilibrium is the best possible outcome for all.²³ This equilibrium is plainly the "correct" coordinating principle, but the tension between its high payoff and its greater riskiness due to strategic uncertainty kept most subjects from choosing the highest effort.

These designs are well suited to testing structural refinements because they involve actions naturally ordered by their payoff implications and labeled accordingly; and the large action spaces and variety of interaction patterns considered allow particularly powerful tests. Applying Harsanyi and Selten's (1988) theory to the stage games in these five treatments predicts 15–52 percent of subjects' initial efforts (Crawford (1991)). Eliminating the priority they give payoff-dominance, allowing risk-dominance (which

embodies most of their ideas about the effects of strategic uncertainty) to determine the predictions in most treatments, yields success rates of 2–52 percent. These results cannot be attributed to the dispersion of subjects' efforts because the theory predicts the modal response in only three of the five treatments (two of five without payoff-dominance). Although there was rapid convergence to equilibrium in four of five treatments, the success rates are no better for last periods: 0–67 percent with and 0–72 percent without payoff-dominance (Crawford (1995a)). Van Huyck, Battalio, and Beil's results are reconsidered from a dynamic point of view in section 6.3.

Contextual principles are also of great importance in coordination.²⁴ In one of the first game experiments, Schelling (1960, pp. 53–67) solicited hypothetical responses to symmetric coordination games in which two players choose among n commonly labeled actions, receiving payoffs of 1 if they choose actions with the same label and 0 otherwise. He focused on contextual principles by combining these games, in which structural principles have no bite, with real action labels such as Heads or Tails, or locations in New York City. The expected payoff of a player who ignores contextual features is $1/n$, independent of his partner's behavior (Crawford and Haller (1990, p. 580)). If, however, players have privately observed personal predilections for labels, whose population frequencies are publicly known, they can normally do better than this by ignoring their own predilections and choosing the label with the highest frequency. If the population frequencies are not a clear guide, they may seek a salient principle that depends only on public knowledge about the labels – a “focal point,” in Schelling's terminology.

Schelling's subjects often exploited their intuitions about how the labels would be used to obtain expected payoffs much greater than $1/n$. Mehta, Starmer, and Sugden (1994) studied this phenomenon in more detail by comparing subjects' action choices when their payoffs did not depend on their own or other subjects' actions with their choices among actions labeled in the same way in coordination games like Schelling's. They interpreted the former choices as personal predilections and the latter as attempts to use the labels to coordinate.

Mehta, Starmer, and Sugden's results for coordination treatments replicated Schelling's, with the frequency of identical choices often several times higher than in the corresponding “personal” treatments. For most sets of labels the population choice frequencies were similarly ordered in both cases, with the popularity of labels in the personal treatment magnified in the coordination treatment, as if subjects were choosing the label with the highest frequency. In some cases the importance of public knowledge was clearly visible. In the personal “Write down any day of the year” treatment, for instance, 88 subjects gave 75 different responses – presumably mostly

“personal” days, but led by December 25 at 5.7 percent. In the corresponding coordination treatment 44.4 percent of the subjects chose December 25; 18.9 percent chose December 10, the day of the experiment; and 8.9 percent chose January 1, all days their public knowledge made more salient than any day their knowledge of predilection frequencies could suggest. Overall, the results provide clear evidence of simultaneous strategic sophistication and the importance of contextual coordinating principles.

5.3 Unstructured bargaining

Some of the most important evidence on simultaneous coordination was provided by a series of bargaining experiments by Roth and his collaborators during the late 1970s and early 1980s (Roth (1987b)). These experiments are of particular interest because they left the bargaining process largely unstructured. This comes closer to bargaining in the field, where rules like those in non-cooperative models of bargaining are seldom encountered. It also allows more informative tests of cooperative and non-cooperative theories of bargaining.

Roth’s designs employed the binary lottery procedure of Roth and Malouf (1979), in which pairs of subjects bargain over a fixed total of 100 lottery tickets, with each subject’s share determining his probability of winning the larger of two possible monetary prizes, specific to him. If subjects could agree on how to share the lottery tickets by an announced deadline the agreement was enforced; otherwise they got zero probabilities. Subjects could make any binding proposal they wished, or accept their partner’s latest proposal, at any time. They could also send non-binding messages at any time, except that they could not identify themselves or, in some treatments, reveal their prizes. The environment was public knowledge, except subjects’ prizes or information about prizes in some treatments.

The designs exploit invariances created by the binary lottery procedure to test both cooperative and non-cooperative theories of bargaining. Under standard assumptions a player maximizes the expected number of lottery tickets he obtains, so that the number of tickets can be taken as his payoff. Cooperative game theory summarizes the implications of a structure by the payoffs players can obtain acting alone or in coalitions. This makes bargaining over a fixed total of lottery tickets equivalent to a complete-information Divide the Dollar game with risk-neutral players, whose symmetry leads cooperative theories to predict equal division of the lottery tickets. These conclusions are independent of players’ risk preferences, prizes, or information about prizes, so that cooperative theories can be tested by observing the effects of varying those factors. Although non-

cooperative theories are harder to test this way because their predictions may depend on the details of the structure, the binary lottery procedure also makes it possible to create invariances that allow such tests, as explained below.

Each treatment paired a subject whose prize was low (typically \$5) with one whose prize was high (typically \$20). A subject always knew his own prize. The first experiment compared two information conditions: “full,” in which a subject also knew his partner’s prize; and “partial,” in which a subject knew only his own prize. The second experiment created a richer set of information conditions using an intermediate commodity, chips, which subjects could later exchange for money in private. A subject always knew his own chip prize and its value in money. There were three information conditions: “high,” in which a subject also knew his partner’s chip prize and its value; “intermediate,” in which a subject knew his partner’s chip prize but not its value; and “low,” in which a subject knew neither his partner’s chip prize nor its value. Subjects were prevented from communicating the missing information, and the information condition was public knowledge.

Partial and low information induce games with identical structures, given that players cannot send messages about chip or money prizes, because their strategy spaces are isomorphic (with chips in the latter treatment playing the role of money in the former) and isomorphic strategy combinations yield identical payoffs (in lottery tickets). For the same reasons full and intermediate information also induce games with identical structures, given that players in the latter cannot send messages about money prizes. Any structural theory, cooperative or non-cooperative, predicts identical outcomes in these pairs of treatments.

A third experiment explored the strategic use of private information by giving subjects the option of communicating missing information about prizes. There were no chips, and a subject always knew his own money prize. There were four basic information conditions: (i) neither subject knew both prizes; (ii) only the subject whose prize was \$20 knew both prizes; (iii) only the subject whose prize was \$5 knew both prizes; and (iv) both subjects knew both prizes. Some treatments made the basic information condition public knowledge, while in others subjects were told only that their partners might or might not know what information they had. Thus there were eight information conditions in all.

I first describe the observed patterns of agreements, and then discuss disagreements. With partial information almost all subjects agreed on a 50–50 division of the lottery tickets. With full information, agreements averaged about halfway between 50–50 and equal expected money winnings, with much higher variance (Roth (1987b, table 2.2)). With low and high information, respectively, agreements averaged close to 50–50 and

roughly halfway between 50–50 and equal expected money winnings, again with higher variance. With intermediate information, agreements averaged close to 50–50 (Roth (1987b, figure 2.1)). Thus partial and low information yielded similar outcomes; but with full and intermediate information, strategically equivalent information about money and chips affected the outcomes in very different ways, which are inconsistent with any structural theory.

The authors attributed the strong influence of subjects' prizes and information about prizes, which are irrelevant in traditional analyses, to the different meanings subjects assigned to chips and money outside the laboratory. Their agreements can be summarized by postulating a commonly understood hierarchy of contextual equal-sharing norms in which subjects implemented the most "relevant" norm their public knowledge allowed, with money most relevant, then lottery tickets, and then chips (Crawford (1990)).²⁵

In the third experiment agreements were largely determined by whether the \$5 subject knew both prizes, clustering around 50–50 when he did not, and shifting more than halfway toward equal expected money winnings when he did (Roth (1987b, table 2.4)). In effect these agreements were determined by the most relevant norm in the above hierarchy that subjects could implement, using their public knowledge plus whatever private information they had incentives to reveal, on the anticipation that it would be used this way. Subjects' revelation decisions were approximately in equilibrium in beliefs in a restricted game in which they could either reveal the truth or nothing at all, when their beliefs are estimated from the mean payoffs in related treatments (Roth (1987b, pp. 27–32)).

There was a subtle interplay between the use of norms and the revelation of private information. In the public-knowledge version of condition (ii) in the third experiment, for instance, the \$5 subject knew that his partner knew which agreement gave them equal expected money winnings, but the \$20 subject usually refused to reveal his prize. This left the 50–50 division the only norm that could be implemented using public knowledge. Although many \$5 subjects voiced suspicions (in transcripts) that they were being treated unfairly, in the end most settled for the 50–50 division. The influence of public knowledge here foreshadowed Mehta, Starmer, and Sugden's (1994) results on contextual focal points.

In all three experiments disagreements occurred, with frequencies ranging from 8–33 percent. Disagreements were most common when both subjects knew enough to implement more than one norm, or when the information condition was not public knowledge.

As explained above, the set of feasible divisions of lottery tickets and subjects' preferences over them were public knowledge, under standard

assumptions, so it is natural to assume complete information in modeling the bargaining game. The non-negligible frequency of disagreements is then incompatible with explanations based on Nash's (1950) bargaining solution or the subgame-perfect equilibrium of an alternating-offers model, as is the strong influence of context on the agreements subjects reached. The manipulation of norms by withholding private information is inconsistent with non-strategic explanations in which subjects "try to be fair." However, most of the results can be understood using a simple strategic model, with players' shared ideas about fairness as coordinating principles.

The model summarizes the strategic possibilities of unstructured bargaining using Nash's (1953) demand game, in which players make simultaneous demands, in this case for lottery tickets. If their demands are feasible they yield a binding agreement; if not there is disagreement. To see how this simple, static game can describe the complex dynamics of unstructured bargaining, assume that delay costs are negligible before the deadline, so that the timing of an agreement is irrelevant. (This is a good approximation for the experiments and many applications to bargaining in the field.) Then, if equilibrium is assumed, all that matters about a player's strategy is the lowest share it can be induced to accept by the deadline. These lowest shares determine the outcome like players' demands in the demand game (Schelling (1960, pp. 267–90), Harsanyi and Selten (1988, pp. 23–4)).

In the complete model, players first decide simultaneously how much private information to reveal. They then bargain, with the ultimate acceptance decisions described by the demand game, in which there is effectively complete information. The demand game has a continuum of efficient equilibria, in which players' demands are just feasible and no worse than disagreement for both. There is also a continuum of inefficient mixed-strategy equilibria with positive probabilities of disagreement. Thus, in this model bargaining is in essence a coordination problem, with players' beliefs the dominant influence on outcomes.

Players' beliefs are focused, if at all, by the most relevant norm their public knowledge (including any revealed private information) allows them to implement. Pure-strategy equilibria, selected this way, yield agreements that closely resemble those observed in the various treatments. From this point of view, it is the desire to avoid a risk of disagreement due to coordination failure that explains \$5 subjects' willingness to settle on the "unfair" 50–50 division in condition (ii) of the third experiment, a phenomenon that is difficult to explain any other way. Finally, mixed-strategy equilibria in which players' beliefs in each treatment are focused on the norms subjects' public knowledge allowed them to implement yield disagreement frequencies close to those observed in the various treatments (Roth (1985)). However, a subsequent, more comprehensive experiment

showed that this model does not fully explain how disagreement frequencies vary with the environment (Roth, Murnighan, and Schoumaker (1988), Roth (1995b, pp. 309–11)).

It is instructive to contrast the view of disagreements as coordination failures suggested by Roth's results with the widespread view that they are explained primarily by asymmetric information about reservation prices. The evidence from the field is equivocal: asymmetric-information bargaining models enjoy some success in explaining strike incidence (Kennan and Wilson (1989)), but there is little evidence that bargaining ceases to be a problem when informational asymmetries are unimportant.

Forsythe, Kennan, and Sopher (1991) conducted an experimental test of a private-information model in which players bargain over the allocation of a "pie" whose size can take two values. One player observes the size and the other knows only its probability distribution; it is common knowledge that disagreement is inefficient in both states; and players can identify some, but not all, of the efficient agreements using common knowledge. With unstructured bargaining there was a non-negligible frequency of disagreements (3–12 percent) even when they were inconsistent with incentive efficiency (Forsythe, Kennan, and Sopher (1991, table 2)). When the pie was small disagreements were more than twice as frequent in treatments in which the informed player could not afford to concede half of the large pie (12–39 percent) than when he could (5–17 percent). Although some of these results are consistent with the incentive-efficiency view of disagreements, they also have a strong flavor of coordination failure.

Once again we find that a complex body of experimental results can be understood by combining traditional equilibrium analysis with empirical knowledge of subjects' ideas about fairness, entering here as coordinating principles rather than payoff perturbations.

6 DYNAMIC EVIDENCE

This section considers evidence on strategic behavior that is dynamic, in that its interpretation depends on how behavior varies over time. Most such evidence has been gathered in environments involving repeated play of a stage game. The typical pattern is an initial period of strategic uncertainty, followed by convergence to an equilibrium in the stage game (in beliefs, if not in actions). Interest usually centers not on convergence, but on how the environment influences the outcome. This influence may depend on complex interactions between the learning dynamics, strategic uncertainty, and the environment, whose effects can persist long after the uncertainty has been eliminated by learning.

Despite this complexity it is often possible to make useful generalizations

about how outcomes are determined. This section discusses the methods and evidence on which such generalizations are based. I begin with evidence from simple environments in which evolutionary game theory yields good predictions of limiting outcomes. I then reconsider some of the “static” evidence from sections 4.2 and 5.2 from a dynamic point of view.²⁶

6.1 Population interactions in simple environments

In simple environments with “evolutionary” structures, the analogy between evolution and learning is often close enough that an evolutionary analysis makes it possible to predict the limiting outcome. Friedman (1996) and Van Huyck, Battalio, and Rankin (1995) studied this issue in abstract two-person 2×2 and 3×3 normal-form games, and Van Huyck *et al.* (1995) studied it in two 3×3 games with structures like Divide the Dollar, one symmetric and one asymmetric.²⁷ The space of aggregate action frequencies was one, two, or (in one case) three dimensional. Except for the restriction to random pairing (or mean matching) the designs address most of the issues about how outcomes are determined in evolutionary games (section 2.3). Friedman’s and Van Huyck, and his collaborator’s results suggest that the aggregate action frequencies often converge to the evolutionary stable outcome whose basin of attraction contains the initial state. This can happen even when that basin of attraction is not the largest one, and equilibrium selection can go against risk-dominance, and/or predictions based on analyses of “long-run equilibria” (section 6.3).

Crawford (1991) studied this issue for Van Huyck, Battalio, and Beil’s (1990, 1991) coordination experiments, finding that the limiting outcomes are surprisingly close to predictions based on evolutionary stability. As discussed in section 6.3, however, because of the larger action spaces and more complex interaction patterns in those experiments a full explanation of the dynamics requires a detailed analysis of learning at the individual level.

6.2 Dominance and iterated dominance revisited

In conjunction with the ultimatum experiments discussed in section 4.2, Prasnikar and Roth (1992) and Roth *et al.* (1991) studied market games, in which nine buyers made offers simultaneously to a single seller, and public-goods games. Although all three games had similar subgame-perfect equilibria, there were large, persistent differences in behavior across treatments, with rapid convergence to the subgame-perfect equilibrium in the public-goods treatment; and nonconvergence, or very slow convergence to a possibly different outcome, in the ultimatum treatment. The authors suggested an informal explanation for these differences based on differences

in out-of-equilibrium payoffs, but their arguments leave room for doubt about whether the payoff differences are large enough to explain the variation in outcomes, or whether the dynamics involve interactions too complex to be understood by “eyeballing” the payoffs.

Roth and Erev (1995) conducted a dynamic analysis of the same data, using a simple model of adaptive learning driven by pecuniary payoffs.²⁸ In their model players choose actions with probabilities determined by “propensities,” which are updated over time according to a formula that yields larger increases for higher realized payoffs. Their adjustment rule satisfies two desiderata from the psychological learning literature, in that the probabilities of actions with higher expected payoffs tend to increase over time (the “law of effect”), but the rate of increase slows over time as players gain experience (the “power law of practice”). Because action choices are random they cannot be viewed as rational responses to beliefs, which are almost always pure for expected-payoff maximizers. However, in stationary environments (and many that are not highly non-stationary) Roth and Erev’s learning rule converges with high probability to a best reply. In this respect it resembles the more sophisticated rules discussed in section 6.3, in which action choices are rational responses to inertial, stochastically convergent beliefs. This resemblance is surprising because Roth and Erev’s rule requires minimal information and is strategically completely unsophisticated: players do not need to know the structure or even that they are playing a game, and they do not need to observe other players’ choices or payoffs.

Roth and Erev investigated the implications of their model by simulation, with the parameters of the learning rule set at the same plausible values for all treatments and initial propensities chosen randomly or estimated from the data for each treatment. The model closely reproduces the dynamics in all three treatments, except that convergence is much slower than in the experiments. Even so, in each case it is the model’s predictions in the intermediate term, not in the long run, that resemble the experimental results. The ultimatum game’s out-of-equilibrium payoffs make the predicted frequencies of low offers by player 1s fall much more quickly than the frequencies of their acceptance by player 2s rises. In all but (possibly) the longest run, this keeps predicted behavior away from the subgame-perfect equilibrium, to an extent that varies with the initial conditions in different countries approximately as in the experiments. The model even reproduces the tendency Roth *et al.* observed for the offers of player 1s to differ increasingly across countries while they converged within each country. By contrast, the very different out-of-equilibrium payoffs in the market and public-goods treatments quickly drive predicted behavior toward the subgame-perfect equilibrium, as in the experiments.

Gale, Binmore, and Samuelson (1995) (inspired by Andreoni and Miller (1995)) conducted a complementary analysis of behavior in ultimatum games based on stochastic replicator dynamics, also driven by pecuniary payoffs. Their model postulates aggregate noise in the dynamics, which interacts with the out-of-equilibrium payoffs in the ultimatum game roughly like the randomness of individual adjustments in Roth and Erev's model. Using a combination of simulation and analytical methods, they showed that if the responses of player 2s near the subgame-perfect equilibrium are noisier than those of player 1s, for whom deviations are much more costly, then even small amounts of noise yield intermediate-term ("long-run" in their terminology) outcomes near an imperfect equilibrium in the ultimatum game that resembles behavior in the experiments.

Nagel (1995) and Stahl (1994) conducted dynamic analyses of the data from Nagel's "guessing game" experiment (section 4.1), in which a subject's sophistication is naturally associated with the number of rounds of iterated dominance implicit in his guesses. Nagel's analysis suggests that sophistication varied across subjects, but had no clear tendency to increase with experience. Stahl postulated more complex learning rules and found evidence of heterogeneity and increasing sophistication over time.

6.3 Simultaneous coordination revisited

Brandts and Holt (1992, 1993a) replicated the results of Banks, Camerer, and Porter's (1994) experiments with signaling games (section 5.1) and conducted new signaling experiments. They found some support for traditional refinements, but they also found considerable strategic uncertainty, which allowed them consistently to obtain convergence to "unrefined" equilibria by varying out-of-equilibrium payoffs. This suggests that outcomes in these games cannot be fully understood without analyzing the learning dynamics.

Van Huyck, Battalio, and Beil (1990, 1991), provide perhaps the clearest evidence on learning and 'history-dependent equilibrium selection. As explained in section 5.2, their subjects played simple coordination games with seven "efforts," in which payoffs were determined by their own efforts and an order statistic of their own and others' efforts. There were five leading treatments, which varied with the order statistic, the number of subjects playing the game, and their interaction pattern. In each case the stage game had seven symmetric, Pareto-ranked equilibria, and a subject's payoff was highest, other things being equal, when his effort equalled the order statistic. In each treatment the stage game was played repeatedly, usually ten times, with the order statistic publicly announced after each play. These environments are a natural setting in which to study the

emergence of conventions to solve coordination problems. Their large action spaces allow rich dynamics, whose variations across treatments discriminate sharply among traditional and different adaptive learning models.

All five treatments had similar initial effort distributions, with high to moderate variances and inefficiently low means, but subjects' subsequent efforts varied across treatments, with persistent consequences for equilibrium selection. In the large-group minimum treatment efforts quickly approached the lowest equilibrium, despite its inefficiency. In the random-pairing minimum treatment efforts slowly approached a moderately inefficient equilibrium, with little or no trend; and in the three median treatments efforts invariably converged to the initial median, although it varied across runs and was usually inefficient. Thus the dynamics were highly sensitive to the size of the groups playing the game and the order statistic, with striking differences in drift, history-dependence, rate of convergence, and the efficiency of the limiting outcome.

Traditional methods do not help to explain these results. Rationality with unrestricted beliefs implies no restrictions on behavior. Equilibrium in the stage game or the repeated game implies some restrictions, but they are the same for every treatment. Predictions based on risk- or payoff-dominance do not reflect the dispersion of initial responses, and differ substantially from subjects' modal initial or final efforts (section 5.2).

Crawford (1995a) and Broseta (1993a and b) proposed adaptive learning models to explain Van Huyck, Battalio, and Beil's results. The models describe players' decisions as rational responses to beliefs, but do not impose equilibrium even in perturbed versions of the game. Instead they use the "evolutionary" structure of the experimental designs to give a flexible characterization of players' learning rules and strategic uncertainty. This permits an informative analysis of the history-dependent learning processes in the experiments, which suggests that the results were strongly influenced by interactions between strategic uncertainty and the learning dynamics. These interactions are not adequately modeled by the mental *tâtonnements* in Harsanyi and Selten's (1988) theory: although perfectly strategically sophisticated players may be able to mentally simulate each other's responses, with strategic uncertainty there is no substitute for analyzing the effects of real feedback.

The specification of learning rules takes advantage of the facts that subjects' payoffs are directly affected by others' efforts only through the order statistic, and that subjects appeared to treat their influences on the order statistic as negligible. On this assumption, their optimal efforts are determined by their beliefs about the current value of the order statistic, so that it suffices to describe the evolution of those beliefs. The model represents

beliefs directly by the optimal efforts they imply, as in the adaptive control literature, rather than as probability distributions or their moments.²⁹ On average each player's beliefs are assumed to adjust linearly part of the way toward the latest observation of the order statistic, in a way that generalizes the fictitious-play and best-reply rules to allow different values of parameters that represent the initial levels, trends, and inertia in beliefs.

Because subjects were externally indistinguishable and had virtually the same information, it does not seem useful to try to explain the differences in their beliefs within the model. Instead the model uses the evolutionary structure of the environment to give a simple statistical characterization of beliefs, in which the average adjustments described above are perturbed each period by idiosyncratic random shocks, which are independently and identically distributed across players, with zero means and given variances.³⁰ These shocks represent strategic uncertainty, described in terms of the differences in players' learning rules. In effect each player has his own theory of coordination, which gives both his initial beliefs and his interpretations of new information an unpredictable component.

Under standard restrictions on the parameters, these learning rules satisfy the law of effect and the power law of practice. They assume less strategic sophistication than a traditional analysis because players ignore their own influences on the order statistic, but more than Roth and Erev's (1995) learning rules because they depend on the best-reply structure. Van Huyck, Battalio, and Beil's subjects seemed to understand the best-reply structure, and it is important to take this into account. Roth (1995a, p. 39, figure 1.2.) found that Roth and Erev's model tracks the dynamics in the large-group minimum treatment much better if it is modified to allow "common learning," in which players' propensities are updated as if they had played the most successful action in the entire population. Because subjects did not usually observe each other's payoffs or actions, the most sensible interpretation of common learning is that players' learning rules incorporated the best-reply structure, and the resulting model yields adjustments close to those of the Crawford and Broseta models.

Specifying the distributions of the shocks yields a Markov process with players' beliefs, represented by their optimal efforts, as the state vector. The transition probabilities may vary over time, as determined by the distributions of the shocks. The dynamics are driven by the dispersion of beliefs, as represented by the variances of the shocks. Different distributional assumptions have different implications for how outcomes are determined, which go a long way toward identifying the stochastic structure.

If the variances of the shocks fall to zero after the first period, so that players differ in their initial beliefs but not in their responses to new observations, the process converges to the equilibrium determined by the

initial realization of the order statistic, independent of the behavioral parameters and the environment. This is consistent with the results in the median treatments, but not with the results in the large-group minimum treatment, where in nine out of nine runs subjects approached an equilibrium below the initial minimum.

If, instead, the variances are positive and remain constant over time, the model is ergodic and allows an analysis of “long-run equilibria” (Kandori (1996)) as in Robles (1997). In the long run the process cycles among the pure-strategy equilibria in the stage game, whose prior probabilities are given by the ergodic distribution. Allowing the variances to approach zero, remaining constant over time, makes the probability of the equilibrium with the lowest (highest) effort approach one for any order statistic below (above) the median. These limits are completely independent of the number of players and the order statistic, as long as the order statistic remains below, or above, the median. (When the order statistic is the median, every pure-strategy equilibrium has positive probability in the limit.) Thus, studying the limiting behavior of an ergodic process with small dispersion leaves most of the questions raised by Van Huyck, Battalio, and Beil’s experiments unanswered.

The dynamics are closest to the experimental results when the variances decline steadily to zero as players learn to forecast the order statistic, as suggested by the power law of practice. If the variances do not decline too slowly the model converges, with probability one, to one of the pure-strategy equilibria of the stage game. Its implications can then be summarized by the prior probability distribution of the limiting equilibrium, which is normally non-degenerate owing to the persistent effects of strategic uncertainty.

The model makes it possible, whether or not the process is ergodic or the dispersion is small, to solve for the entire history of players’ beliefs and efforts as functions of the behavioral parameters, the shocks, the number of players, and the order statistic. The outcome is built up period by period from the shocks, whose effects persist indefinitely. This persistence makes the process resemble a random walk, in the aggregate, but with possibly non-zero drift that depends on the behavioral parameters, the variances, and the environment; and declining variances that allow the process to converge to a particular equilibrium. This limiting equilibrium is normally sensitive to the entire history of players’ interactions.

The model allows a comparative dynamics analysis, which shows both qualitatively and quantitatively how strategic uncertainty interacts with the environment (holding the behavioral parameters constant across treatments) to determine the outcome. The quantitative analysis is based in part on analytical approximations of the drift of the process. These reveal

that in the median and random-pairing minimum treatments the drift is zero, and that in the large-group minimum treatment the drift is increasingly negative with larger numbers of players, and proportional to the standard deviation that represents the dispersion of beliefs. These results and analogous approximations of the variances suggest patterns of variation across treatments like those in the experiments.³¹

To develop the model's full implications, or to test it, the behavioral parameters and the variances that represent strategic uncertainty must be evaluated. The model makes it possible to estimate the parameters econometrically, using the data from each treatment. The estimated parameters satisfy the restrictions suggested by the theory. The variances that represent the dispersion of beliefs are initially large and decline gradually to zero, and the hypothesis that subjects had identical beliefs throughout is strongly rejected. Using repeated simulation to infer the estimated model's implications confirms the accuracy of the approximations, and shows that the model provides an adequate statistical summary of subjects' behavior while reproducing the dynamics of their interactions in all five treatments. In this sense, taking the effects of strategic uncertainty into account yields a simple, unified explanation of Van Huyck, Battalio, and Beil's results.

Crawford and Broseta (1995) proposed a similar model to explain the results of Van Huyck, Battalio, and Beil's (1993) experiment, which modified one of the nine-player median treatments from their 1991 experiment by auctioning the right to play the same nine-person median game each period in a group of 18. The winners were charged the same market-clearing price, which was publicly announced each period before they played the median game. The auctions can be expected to enhance efficiency because subjects' beliefs usually differ, auctions select the most optimistic subjects, and the game is one in which optimism favors efficiency. The subjects did much better than this argument suggests, quickly bidding the price up to a level that could be recouped only in the most efficient equilibrium and then converging to that equilibrium. The dynamics focused their beliefs as in the intuition for forward induction refinements, in which players infer from other players' willingness to pay to play a game that they expect payoffs that repay their costs, and will play accordingly. This suggests an important new way in which competition may foster efficiency.

Surprisingly, Crawford's (1995a) and Broseta's (1993a and b) methods can be adapted to analyze the dynamics in this more complex environment. The results show how the strength of the efficiency-enhancing effect of auctions is determined by the environment and the behavioral parameters, apportioning it among an order-statistic effect like the one that drives the dynamics in the earlier models, modified by the "optimistic subjects" effect

just described and a “forward induction” effect like the one just described. The estimated model suggests that these effects contributed roughly equally to the efficiency-enhancing effect of auctions in Van Huyck, Battalio, and Beil’s environment, and that auctions will have similar but possibly weaker effects in nearby environments with different numbers of players, different order statistics, and different degrees of competition for the right to play.

These analyses suggest that it will often be possible to analyze the history-dependent learning processes commonly observed in experiments. Once again the models suggested by the experimental results are hybrids, in these cases combining the “evolutionary” structure of the experimental designs with simple characterizations of individual learning, with empirical parameters that reflect the structure of learning rules, the initial level of strategic uncertainty, and the rate at which it is eliminated by learning.

7 CONCLUSION

This chapter has surveyed a large body of experimental work with well thought-out designs and the careful control needed to test strategic models, which addresses issues central to the analysis of strategic interaction. I hope that my discussion conveys some of the richness of the possibilities of experiments, and gives some indication of the extent to which thinking about their results can suggest fruitful new directions for theoretical work.

Although the laboratory is not the field, many experimental results are so robust and so coherent that it is difficult to dismiss them as unrepresentative of “real” behavior. The notion that behavior is a rational response to beliefs, in conjunction with ideas from traditional non-cooperative and cooperative game theory, evolutionary game theory, and adaptive learning models, is surprisingly helpful in organizing the data. In no way, however, do the results justify the traditional view that rationality is all that is needed to understand strategic behavior. Most subjects seem to have some strategic sophistication, but seldom enough to justify an analysis based exclusively on equilibrium, however refined. Moreover, what sophistication they have often takes non-traditional forms, and their beliefs are more likely to be coordinated by inductive and/or contextual principles than deductive and/or structural ones. When subjects’ beliefs are not coordinated at the start, learning commonly yields convergence to an equilibrium in the stage game; but the outcome is frequently history-dependent, and the effects of strategic uncertainty may persist long after it has been eliminated by learning. In such cases both traditional refinements and overly simple models of adaptive learning or evolutionary dynamics may predict poorly.

Nonetheless, the results of experiments give good reason to hope that

most strategic behavior can be understood via a synthesis that combines elements from each of the leading theoretical frameworks with a modicum of empirical information about behavior, in proportions that vary with the environment in predictable ways. In this synthesis theory will play a wider role than in most strategic analyses to date, providing a framework within which to learn which ideas are useful and which aspects of behavior cannot reliably be determined by theory, and to gather the empirical information needed to close the model.

The analysis of such models will require new static methods that combine rationality with empirically sensible restrictions on strategies, without imposing coordination of beliefs, as in Cho (1994), Rabin (1993, 1994), and Watson (1993, 1996). It will also require new dynamic methods that take the persistent effects of strategic uncertainty in history-dependent learning processes fully into account, and that go beyond random pairing to consider other interaction patterns that are important in economics, as in Roth and Erev (1995), Crawford (1995a), Broseta (1993a and b), and Crawford and Broseta (1995).

Notes

Invited Symposium on Experimental Economics, Econometric Society Seventh World Congress, Tokyo, August 1995. I owe thanks to John McMillan, Alvin Roth, Joel Sobel, and especially Mark Machina for helpful advice and to Miguel Costa Gomes for able research assistance. My debt to Thomas Schelling and the many experimentalists and theorists who have since studied behavior in games should be clear from the text.

- 1 There is nonetheless a history of valuable empirical work using field data from strategic environments, usually with well-specified, readily observable structures.
- 2 See also Plott (1991), Roth (1987a, 1991, 1995a), and Smith (1989).
- 3 This focus and space limitations have led me to exclude a great deal of important experimental work; see Kagel and Roth (1995) for a comprehensive survey.
- 4 Hey's (1996) companion chapter surveys experiments on individual decisions. Conlisk (1996) gives a good overview of bounded rationality in decisions and games.
- 5 Compare the notion of "theory of mind" in cognitive psychology, where it has been found experimentally that some aspects of what I call strategic sophistication develop in normal (but not autistic) children around age three (Leslie (1994)).
- 6 Dekel and Gul (1996) give a good overview of this approach. Common knowledge of beliefs can be relaxed to approximate common knowledge for strict equilibria, and to mutual knowledge for two-person games.
- 7 Beliefs could also be coordinated by preplay communication, but communication does not always yield equilibrium in the underlying game (Aumann (1990)).
- 8 Making this distinction precise is difficult and probably premature at this stage of our empirical knowledge. For instance, a game theorist might find the analogies

created by Nash's (1950) axioms so transparent that applying his bargaining solution is just copying behavior inductively from one problem to another, but others are likely to view the Nash solution as a deductive theory (Schelling (1960, pp. 113–14)).

- 9 With random pairing stable frequencies are also in equilibrium in the game played by pairs. Some qualifications apply for finite populations or extensive-form stage games.
- 10 Individual pairs can of course play asymmetric action combinations by chance, but asymmetric aggregate frequencies are statistically unplayable, even in the limit. Crawford and Haller (1990, p. 580) give a "traditional" analog of this argument. Evolutionary game theory also has a way to model the effects of distinguished roles, illustrated by the Stoplight example of Crawford (1991, section 3).
- 11 The founding analysis of evolutionary game theory, Fisher's (1930) explanation of the tendency of the ratio of male to female births to remain near 1, is a game against the field (Maynard Smith (1982, pp. 23–7)). This problem – one of the most beautiful in science – requires a game-theoretic explanation because a ratio of 1 equalizes the fitnesses of having male and female offspring, and does not maximize the growth rate of the entire population. The model is a game against the field because the fitnesses depend (nonlinearly) on the population frequencies of male and female offspring.
- 12 By contrast, "rational" learning models such as Bray and Kreps (1987) and Crawford and Haller (1990) assume equilibrium in the repeated game.
- 13 In models that describe action choices directly, without reference to beliefs, I use "strategic uncertainty" loosely to refer to systematic deviations from equilibrium.
- 14 Preplay communication and repeated-game strategies, which raise similar analytical issues but deserve separate treatment because of their economic importance, are discussed in Crawford (1995b, sections 7 and 8), omitted here due to space limitations.
- 15 "Large" populations are surprisingly small: subjects usually treat individual influences as negligible in groups of 10–15, and sometimes in groups as small as five.
- 16 There is also a large body of experiments on nonpecuniary effects (see, for example, Camerer and Thaler (1995) and Roth (1995b)).
- 17 A survey in progress should reveal if failure to rely on dominance is more prevalent among those who were taught to look both ways before crossing one-way streets.
- 18 Thus it gathers the same information as eye-movement studies of problem-solving in psychology, but more systematically. Subjects were not allowed to record their payoffs, and their look-up patterns suggest that they did not memorize them.
- 19 Less time is spent checking the first-period payoff because as long as it is higher than the second-period payoff it does not affect the subgame-perfect equilibrium offer.

- 20 If the pie is discrete, there is also a nearby subgame-perfect equilibrium in which player 1 offers player 2 the smallest feasible positive amount and player 2 accepts.
- 21 Thus the modal offers are consistent with equilibrium in beliefs. Equilibrium would also require optimality of the other offers and consistency of the rejections of player 2s with the hypothesized distribution of nonpecuniary payoffs discussed below.
- 22 In coordination, for instance, it can be just as disadvantageous to be “too clever” for one’s partners as to be not clever enough – provided that one’s cleverness does not include the ability to predict the effects of others’ lack of cleverness.
- 23 If players are risk neutral these conclusions extend to the stage game that describes the entire population’s interactions in the random-pairing minimum treatment, with its median effort as the order statistic (Crawford (1995a, fn. 10, p. 110)).
- 24 The nun who taught me in third grade that Jesus was *exactly* six feet tall had an intuitive grasp of the importance of contextual principles, if not of their transience.
- 25 The equal-chip and equal-money norms are contextual because they depend on things that do not affect the feasible divisions of lottery tickets or subjects’ preferences over them. Bar-Hillel and Yaari (1993) surveyed students’ views about some less abstract norms, in which needs and beliefs enter as well as preferences.
- 26 Crawford (1995b, section 6.2), discusses analyses that explain “cooperative” deviations from subgame-perfect or sequential equilibrium in experiments with multi-period extensive-form games like the Centipede Game and the finitely repeated Prisoners’ Dilemma by hypothesizing idiosyncratic, privately observed payoff or action perturbations. Crawford (1995b, section 6.4), discusses the small body of evidence on the structure of learning rules. Both topics are omitted here owing to space limitations.
- 27 Friedman often speeded convergence using *mean matching*, in which the payoffs of random pairing are simulated without the uncertainty of matching in a game against the field, by replacing them by their expectations given the population frequencies.
- 28 The complexity of dynamic models makes it natural to start by considering only pecuniary payoffs. However, the inertia of Roth and Erev’s (1995) adjustment process and their estimation of subjects’ initial responses from the data yield player 2s behavior similar to the static model with nonpecuniary payoffs proposed in section 4.2.
- 29 Players’ efforts are actually determined by a discrete-choice model, in which their beliefs are the continuous latent variables. I ignore this distinction in the text.
- 30 Broseta (1993a and b) obtains similar results for a stochastic structure with richer dynamics.
- 31 The analysis shows that it was no coincidence that the most interesting dynamics were found in a game against the field, the large-group minimum

treatment: random pairing eliminates the effects of strategic uncertainty that drove those results.

References

- Andreoni, James and Miller, John (1995). "Auctions with artificial adaptive agents." *Games and Economic Behavior*, 10: 39–64.
- Aumann, Robert (1990). "Nash equilibria are not self-enforcing." In Gabszewicz, J. J. Richard, J.-F. and Wolsey, L. A. (eds.), *Economic Decision-Making: Games, Econometrics and Optimization*. Lausanne: Elsevier Science Publishers, pp. 201–6.
- Aumann, Robert and Brandenburger, Adam (1995). "Epistemic conditions for Nash equilibrium." *Econometrica*, 63: 1161–80.
- Banks, Jeffrey, Camerer, Colin, and Porter, David (1994). "An experimental analysis of Nash refinements in signaling games." *Games and Economic Behavior*, 6: 1–31.
- Bar-Hillel, Maya and Yaari, Menahem (1993). "Judgments of distributive justice." In Mellers, Barbara and Baron, Jonathan (eds.), *Psychological Perspectives on Justice: Theory and Applications*. New York: Cambridge University Press.
- Beard, T. Randolph and Beil, Richard (1994). "Do people rely on the self-interested maximization of others?: an experimental text." *Management Science*, 40: 252–62.
- Brandts, Jordi and Holt, Charles (1992). "An experimental test of equilibrium dominance in signaling games." *American Economic Review*, 82: 1350–65.
- (1993a). "Adjustment patterns and equilibrium selection in experimental signaling games." *International Journal of Game Theory*, 22: 279–302.
- (1993b). "Dominance and forward induction: experimental evidence." In Isaac R. Mark (ed.), *Research in Experimental Economics*, Vol. V. Greenwich, CT: JAI Press, pp. 119–36.
- Bray, Margaret and Kreps, David (1987). "Rational learning and rational expectations." In Feiwel, George (ed.), *Arrow and the Ascent of Modern Economic Theory*. New York: New York University Press, pp. 597–625.
- Broseta, Bruno (1993a). "Strategic uncertainty and learning in coordination games." UCSD Discussion Paper 93–34.
- (1993b). "Estimation of a game-theoretic model of learning: an autoregressive conditional heteroskedasticity approach." UCSD Discussion Paper 93–35.
- Camerer, Colin, Johnson, Eric, Rymon, Talia, and Sen, Sankar (1993). "Cognition and framing in sequential bargaining for gains and losses." In Binmore, Kenneth, Kirman, Alan, and Tani, Piero (eds.), *Frontiers of Game Theory*. Cambridge, MA: MIT Press, pp. 27–47.
- Camerer, Colin and Thaler, Richard (1995). "Anomalies: ultimatums, dictators and manners." *Journal of Economic Perspective*, 9: 209–19.
- Cho, In-Koo (1994). "Stationarity, rationalizability and bargaining." *Review of Economic Studies*, 61: 357–74.
- Conlisk, John (1996). "Why bounded rationality?" *Journal of Economic Literature*, 34: 669–700.

- Cooper, Russell, DeJong, Douglas, Forsythe, Robert, and Ross, Thomas (1994). "Alternative institutions for resolving coordination problems: experimental evidence on forward induction and preplay communication." In Friedman, James (ed.), *Problems of Coordination in Economic Activity*. Boston: Kluwer, pp. 129–46.
- Crawford, Vincent (1989). "Learning and mixed-strategy equilibria in evolutionary games." *Journal of Theoretical Biology*, 140: 537–50. Reprinted in Bicchieri, Chrislina, Jeffrey, Richard, and Skyrms, Brian (eds.), *The Dynamics of Norms*. New York: Cambridge University Press.
- (1990). "Explicit communication and bargaining outcomes." *American Economic Review Papers and Proceedings*, 80: 213–19.
- (1991). "An 'evolutionary' interpretation of Van Huyck, Battalio, and Beil's experimental results on coordination." *Games and Economic Behavior*, 3: 25–59.
- (1995a). "Adaptive dynamics in coordination games." *Econometrica*, 63: 103–43.
- (1995b). "Theory and experiment in the analysis of strategic interaction." UCSD Discussion Paper 95–37.
- Crawford, Vincent and Broseta, Bruno (1995). "What price coordination? Auctioning the right to play as a form of preplay communication." UCSD Discussion Paper 95–41.
- Crawford, Vincent and Haller, Hans (1990). "Learning how to cooperate: optimal play in repeated coordination games." *Econometrica*, 58: 571–95.
- Dekel, Eddie, and Gul, Faruk (1996). "Rationality and knowledge in game theory." Chapter 5 in this volume.
- Fisher, R. A. (1930). *The Genetical Theory of Natural Selection*. Oxford: Clarendon Press.
- Forsythe, Robert, Kennan, John and Sopher, Barry (1991). "An experimental analysis of strikes in bargaining games with one-sided private information." *American Economic Review*, 81: 253–70.
- Friedman, Daniel (1991). "Evolutionary games in economics." *Econometrica*, 59: 637–66.
- (1996). "Equilibrium in evolutionary games: some experimental results." *Economic Journal* 106: 1–25.
- Fudenberg, Drew and Kreps, David (1993). "Learning mixed equilibria." *Games and Economic Behavior*, 5: 320–67.
- Gale, John, Binmore, Kenneth, and Samuelson, Larry (1995). "Learning to be imperfect: the ultimatum game." *Games and Economic Behavior*, 8: 56–90.
- Harsanyi, John and Selten, Reinhard (1988). *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.
- Hey, John (1996). "Experiments and the economics of individual decision making under risk and uncertainty." Chapter 6 in this volume.
- Kagel, John and Roth, Alvin (eds.) (1995). *Handbook of Experimental Economics*. Princeton: Princeton University Press.
- Kandori, Michihiro (1996). "Evolutionary game theory in economics." Chapter 8 in this volume.
- Kennan, John and Wilson, Robert (1989). "Strategic bargaining models and

- interpretation of strike data." *Journal of Applied Econometrics*, 4 (Supplement): S87–S130.
- Leslie, Alan M. (1994). "Pretending and believing – issues in the theory of TOMM." *Cognition*, 50: 211–38.
- Marimon, Ramon (1996). "Learning from learning in economics." Chapter 8 in this volume.
- Maynard Smith, John (1982). *Evolution and the Theory of Games*. New York: Cambridge University Press.
- Mehta, Judith, Starmar, Chris, and Sugden, Robert (1994). "The nature of salience: an experimental investigation of pure coordination games." *American Economic Review*, 84: 658–73.
- Nagel, Rosemarie (1995). "Unraveling in guessing games: an experimental study." *American Economic Review*, 85: 1313–26.
- Nash, John (1950). "The bargaining problem." *Econometrica*, 18: 155–62.
(1953). "Two-person cooperative games." *Econometrica*, 21: 128–40.
- Plott, Charles (1991). "Economics in 2090: the views of an experimentalist." *Economic Journal*, 101: 88–93.
- Prasnikar, Vesna and Roth, Alvin (1992). "Considerations of fairness and strategy: experimental data from sequential games." *Quarterly Journal of Economics*, 107: 865–88.
- Rabin, Matthew (1993). "Incorporating fairness into game theory and economics." *American Economic Review*, 83: 1281–302.
(1994). "Incorporating behavioral assumptions into game theory." In Friedman, James (ed.), *Problems of Coordination in Economic Activity*. Boston: Kluwer, pp. 69–86.
- Robles, Jack (1997). "Evolution and long run equilibria in coordination games with summary statistic payoff technologies." *Journal of Economic Theory*, forthcoming.
- Roth, Alvin (1985). "Toward a focal-point theory of bargaining." In Roth, Alvin (ed.), *Game-Theoretic Models of Bargaining*. New York: Cambridge University Press, pp. 259–68.
(1987a). "Laboratory experimentation in economics." In Bewley, Truman (ed.), *Advances in Economic Theory: Fifth World Congress*, Econometric Society Monograph No. 12. New York: Cambridge University Press, pp. 269–99.
(1987b). "Bargaining phenomena and bargaining theory." In Roth, Alvin (ed.), *Laboratory Experimentation in Economics: Six Points of View*. New York: Cambridge University Press. pp. 14–41.
(1991). "Game theory as a part of empirical economics." *Economic Journal*, 101: 107–14.
(1994). "Let's keep the con out of experimental econ." *Empirical Economics*, 19: 279–89.
(1995a). "Introduction to experimental economics." In Kagel, John and Roth, Alvin (eds.), pp. 3–109.
(1995b). "Bargaining experiments." In Kagel, John and Roth, Alvin (eds.), pp. 253–348.

- Roth, Alvin and Erev, Ido (1995). "Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term." *Games and Economic Behavior*, 8: 164–212.
- Roth, Alvin and Malouf, Michael. (1979). "Game-theoretic models and the role of information in bargaining." *Psychological Review*, 86: 574–94.
- Roth, Alvin, Murnighan, J. Keith and Schoumaker, Françoise (1988). "The deadline effect in bargaining: some experimental evidence." *American Economic Review*, 78: 806–23.
- Roth, Alvin, Prasnikar, Vesna, Okuno-Fujiwara, Masahiro, and Zamir, Shmuel (1991). "Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: an experimental study." *American Economic Review*, 81: 1068–95.
- Samuelson, Larry (1993). "Does evolution eliminate dominated strategies?" In Binmore, Kenneth, Kirman, Alan, and Tani, Piero (eds.), *Frontiers of Game Theory*. Cambridge, MA: MIT Press, pp. 213–35.
- Schelling, Thomas (1960). *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- (1978). *Micromotives and Macrobehavior*. New York: W.W. Norton.
- Smith, Vernon (1989). "Theory, experiment and economics." *Journal of Economic Perspectives*, 3: 151–69.
- Stahl, Dale (1994). "Rule learning in a guessing game." Working Paper 9415, University of Texas.
- Stahl, Dale and Wilson, Paul (1995). "On players' models of other players: theory and experimental evidence." *Games and Economic Behavior*, 10: 218–54.
- Van Huyck, John, Battalio, Raymond, and Beil, Richard (1990). "Tacit coordination games, strategic uncertainty, and coordination failure." *American Economic Review*, 80: 234–48.
- (1991). "Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games." *Quarterly Journal of Economics*, 106: 885–910.
- (1993). "Asset markets as an equilibrium selection mechanism: coordination failure, game form auctions, and tacit communication." *Games and Economic Behavior*, 5: 485–504.
- Van Huyck, John, Battalio, Raymond, Mathur, Sondip, Ortmann, Andreas, and Van Huyck, Patsy (1995). "On the origin of convention: evidence from symmetric bargaining games." *International Journal of Game Theory*, 24: 187–212.
- Van Huyck, John, Battalio, Raymond and Rankin, Frederick (1995). "On the origin of convention: evidence from coordination games." Research Report 4, Texas A&M University Economics Laboratory.
- Watson, Joel (1993). "A reputation refinement without equilibrium." *Econometrica*, 61: 199–205.
- (1996). "Reputation in repeated games with no discounting." Forthcoming in *Games and Economic Behavior*.