

Inferring repeated-game strategies from actions: evidence from trust game experiments[★]

Jim Engle-Warnick¹ and Robert L. Slonim²

¹ Department of Economics, McGill University, Montreal, Quebec, H3A 2T7, CANADA
(e-mail: jim.enge-warnick@mcgill.ca)

² Department of Economics, Weatherhead School of Management, Case Western Reserve University,
Cleveland, OH 44106, USA (e-mail: robert.slonim@case.edu)

Received: December 23, 2002; revised version: April 19, 2005

Summary. Combining a strategy model, an inference procedure and a new experimental design, we map sequences of observed actions in repeated games to unobserved strategies that reflect decision-makers' plans. We demonstrate the method by studying two institutional settings with distinct theoretical predictions. We find that almost all strategies inferred are best responses to one of the inferred strategies of other players, and in one of the settings almost all of the inferred strategies, which include triggers to punish non-cooperators, are consistent with equilibrium strategies. By developing a method to infer unobserved repeated-game strategies from actions, we take a step toward making game theory a more applied tool, bridging a gap between theory and observed behavior.

Keywords and Phrases: Repeated games, Strategies, Finite automata, Trust, Experimental economics.

JEL Classification Numbers: C72, C80, C90.

* The authors are indebted for discussions with Ray Battalio, David Cooper, Robin Dubin, John Duffy, Ellen Garbarino, Susan Helper, Margaret Meyer, John Miller, Jim Rebitzer, Mari Rege, Al Roth, and John Van Huyck. The authors also benefited from discussions at economic department seminars at Case Western Reserve, McMaster and McGill University, University of Pittsburgh, SUNY-Stony Brook, and Texas A&M, and participants at the 2002 European Winter Meeting of the Econometric Society. We are grateful for the financial support provided by the Department of Economics at the University of Pittsburgh and Case Western Reserve University.

Correspondence to: R.L. Slonim

1 Introduction

Although there is extensive theoretical analysis of repeated games, surprisingly little evidence exists regarding the strategies that a heterogeneous population actually uses when playing them. Knowing the strategies people use is important for several reasons. First, many economic transactions occur in the context of institutions that involve repeated interactions. Second, researchers will be able to go beyond explaining *what* happened in repeated games, e.g., in carefully designed experiments, to explaining *why* it happened. Third, empirically supported hypotheses regarding strategies provide insights into behavior that may feed back into theories: our results will show that relatively few strategies explain the vast majority of behavior, providing empirical evidence for equilibrium selection. Thus, an empirically based model of repeated-game strategies opens areas of research to better understand behaviour across a variety of institutions involving repeated interactions.

The need to infer strategies from actions exists because strategy choices are not observed, and cannot be directly recovered, in most settings. For example, we can observe central bank actions (raise, maintain or lower the short-term interest rate), union actions (go on strike or don't go on strike) or a firm's actions (increase or decrease orders), but we cannot observe the strategies that generate these actions. In this paper we develop a method to infer repeated-game strategies from actions without interfering with the decision-making process. The method consists of a strategy model, an inference procedure and a new experimental design.

Our approach searches for deterministic repeated-game strategies that reproduce the actions observed in experimental choice data. We allow for the possibility of inferring multiple strategies, thus allowing for heterogeneity across players, and for individual players mixing over their strategies. We also do not attempt to replicate all of the observed strategy choices, thus allowing for trembles or seldomly-used strategies.

In making choices regarding modeling repeated-game strategies, the experimental design and the inference procedure, the most important task is to reduce the dimensionality of the problem of identifying unobserved strategy choices from the observed actions of decision makers. The strategy model we examine uses finite automata, which only include a subset of the theoretically large set of possible strategies in repeated games. To further reduce the number of strategies we consider, we designed an experiment to minimize the number of contingencies that might influence the decision maker; specifically, we examine a stage game with only two players, each with two actions, and we examine a repeated game consisting of, on average, only five repetitions of the stage game. And to reduce the number of strategies we infer while allowing for heterogeneity in strategy choices (e.g., across time and across the population), we impose a cost proportional to the amount of heterogeneity (i.e., number of strategies) that we infer.

We apply our method to examine play in a repeated trust game. In the trust stage game, two players move sequentially. The first player may trust by relinquishing ownership of an endowment to the second player. If the first player trusts, then the second player receives the first player's endowment plus a surplus and must decide whether to return part of the gains or to keep everything. In finitely repeated games

(henceforth finite games), two players play the stage game exactly five times. In indefinitely repeated games (henceforth indefinite games), two players play each stage game knowing that there is a chance that the stage game will be the last they will play together. We examine the inference method across these institutions since they have distinct equilibrium strategies, thus allowing us to test the robustness of the method.

Using the strategy inference method, in the indefinite games we find that subjects play trigger strategies that are consistent with equilibrium strategies, and in the finite games we find that subjects play strategies with end-game effects. In the finite games we also find that the strategies we infer evolve over time in a manner consistent with unravelling behavior. And in both the finite and indefinite games we find that almost all of the inferred strategies are best responses to one of the inferred strategies of the opponents. Thus, the analysis shows the ability of the methodology to find a subset of strategies both on and off the equilibrium path that are theoretically justified.

The direct inference methodology introduced in this paper, in which strategies are inferred directly from actions, is complementary to several existing approaches of strategy inference. In the “strategy method” of Selten, Mitzkewitz, and Ulrich (1997), strategy choices are made observable through direct elicitation; a second approach is to validate inferred strategies by tracking the manner in which subjects collect and process information (i.e., to collect attentional data as in Costa-Gomes et al. 2001; Johnson et al. 2002); a third approach is to estimate a probabilistic choice model as in El-Gamal and Grether (1995), Engle-Warnick (2003), Selten and Stoeker (1986), and Stahl and Wilson (1995); other approaches involve experimental manipulation and protocol responses.¹

We contribute to the literature by studying *repeated-game* strategies and by studying an environment in the laboratory that we believe resembles decision-making in the field.² Specifically, we do not require subjects to explicitly consider every contingency (the strategy method), nor do we introduce a layer of decision-making in front of the subjects’ choices of actions (the information tracking method). Thus, our approach minimizes the interaction between the methodology and the decision-maker to more closely resemble decision-making in the field.

The remainder of the paper proceeds as follows. Section 2 presents the experimental design and provides a preliminary discussion of the results. The following three sections present the strategy inference method: Section 3 presents the strategy model, Section 4 presents the inference procedure, and Section 5 presents the strategies inferred. Section 6 discusses the advantage of the approach and insights gained. The conclusion offers directions for future research.

¹ For evidence on how these approaches may differ, see comparisons of the strategy and direct methods of inference by Slonim (1994), and Brandts and Charness (2000).

² The closest related study of repeated games is the repeated PD game of Selten and Stoeker (1986). Van Huyck et al. (2001) study indefinitely repeated games using a similar stage game, with a continuum of actions.




Figure 1. The trust stage game

2 The trust game experiment

As described throughout this section, and the remainder of the paper, the experimental design is integral to determining what strategies people play in repeated games. Figure 1 shows the extensive form game we study experimentally.³ Two players, Player A and B, are each given an endowment of \$0.40 to start each stage game. Player A chooses between the action *Send* (*S*) and *Don't Send* (*D*). If Player A chooses *D*, then both players receive their endowment and the stage game ends. If Player A chooses *S*, then Player A's endowment is doubled and given to Player B (e.g., reflecting a return on an investment). Player B then chooses between the action *Keep* (*K*) and *Return* (*R*). If Player B chooses *K*, then Player B receives \$1.20 and Player A receives \$0.00. If Player B chooses *R*, then both players receive \$0.60.

2.1 Experimental procedures

Subjects were randomly assigned to be Player A or B for an entire session. At the beginning of each supergame every Player A was randomly and anonymously paired with a Player B. Each session began with twenty indefinite supergames and finished with twenty finite supergames.⁴ We chose to examine indefinite and finite supergames since theory predicts, and past evidence and intuition suggests, the two institutions may induce different strategy choices.

To play indefinite supergames, subjects were told that at the end of each round (i.e., stage game) within every supergame there was a continuation probability of $p = 0.8$. If the supergame continued, then subjects would play another round with the same opponent. If the supergame ended, then subjects would be randomly and anonymously paired with a new opponent to begin a new supergame. We randomly drew sequences of supergame lengths (i.e., the number of rounds in each supergame) prior to running the sessions. The average length per supergame was 5.1; the longest was fourteen rounds, and the shortest was one round. We ran two sessions with each sequence of supergame lengths and used one for the training sample and the other

³ This game is theoretically analyzed by Kreps (1990). Berg et al. (1995) examine behavior when a similar game is played without repetition.

⁴ Instructions are available from the authors.

for the test sample. Having the same sequence of lengths in the training and test samples eliminates a potential source of variation between them.

After completing twenty indefinite supergames subjects were told that all future supergames would last exactly five rounds. Twenty additional supergames were played under these finitely repeated conditions. Similar to the indefinite treatment, subjects were anonymously and randomly paired with a new opponent to begin each finite supergame. We chose the finite supergames to last five rounds so that the number of rounds in the finite condition would equal the expected number of rounds of the indefinite supergames.

We chose the expected length of supergames to be five rounds to make strategy inference more plausible. If supergame lengths are much shorter than five rounds (e.g., one or two rounds), then there won't be much opportunity to observe how players react to opponent choices and thus identification among different strategies will not be possible. On the other hand, with longer supergame lengths (e.g., twenty or thirty rounds), the possibility of players changing strategies or making mistakes during the course of the supergame increases. Also, the longer the length of each supergame, the fewer the number of supergames we can observe. An average length of five rounds appealed to us as long enough to observe players reacting to opponent actions while short enough to minimize time-variance concerns.⁵

Four sessions were run with fourteen subjects in each. There were thus 560 supergame observations per role in the indefinite condition (seven of each player type times four sessions times twenty supergames) and an additional 560 supergame observations per role in the finite condition.⁶ The experiments were run at the University of Pittsburgh. Subjects were paid a \$5.00 participation fee plus their earnings from four supergames that were randomly selected at the end of the session.

2.2 *Trust game equilibria*

The unique subgame perfect equilibrium of the stage game is for Player B to play K if Player A plays S and for Player A to thus play D . Backward induction leads to the same subgame perfect equilibrium behavior for both players in every round of the finite game. For the indefinite game there are many equilibria.⁷ At one extreme, players may play S and R every round and at the other extreme Player A may play D in every round. Modeling repeated-game strategies with finite automata does not

⁵ The experimental design represents a trade-off between inducing as many game histories as possible and allowing subjects to be paired with the same player multiple times. However, subjects never knew who they were currently playing against. Nor did they know if or when they played against the current opponent in the past, and if or when they would play against the current opponent in the future. The issue of possibly encountering an opponent in a later supergame was not mentioned by a single subject in post-session protocols that asked subjects to describe their strategic decision-making. Our view is that there is no way to avoid this trade-off when attempting to infer repeated-game strategies; because there are so many possible game histories, we chose to induce as many as possible, while trying to control for independence to the extent possible, both with the experimental design and the inference procedure.

⁶ Due to a computer failure we lost seven indefinite observations and 98 finite observations in one session. For this reason there are 553 indefinite and 462 finite observations. We chose the session that experienced the computer failure to be one of the two sessions for the test sample.

⁷ This is a result from the Folk Theorem of Repeated Games, see Fudenberg and Maskin (1986).

change the characterization of equilibrium when strategy complexity is not taken into account (see Abreu and Rubinstein, 1988).

The continuation probability was chosen to induce actions to vary. If the discount factor (i.e., the continuation probability) is less than 0.75, then a risk neutral Player B's expected payoff is maximized by always playing *K*, regardless of Player A's strategy. We thus chose a discount factor greater than 0.75 to permit cooperative equilibria. However, we also chose the discount factor to be near the 0.75 to make cooperation difficult to achieve. Two paper and pencil pilot sessions confirmed that the continuation probability of 0.80 induced considerable variation in actions across supergames.⁸

2.3 Preliminary results

Our experiments were designed so that we can interpret sequences of actions taken by a subject within a supergame as having been generated by a stationary and unobservable repeated-game strategy. Since repeated-game strategies map game histories into actions, we begin by investigating game histories that may figure prominently in subject decision making. Specifically, we select particular game histories that are either relevant to theoretical predictions or due to past experimental evidence by introducing an indicator variable that takes on a value of one whenever a particular history has occurred within a supergame and a value of zero whenever it has not.

We estimate the effect of the history on the probability of taking the actions *S* and *R* using a fixed effects logit models. To more directly examine how past supergame lengths affected choice, we estimate the probability that Trustor *i* (Trustee *i*) played $S_{i,r}$ ($R_{i,r}$) in Round *r* using the following two fixed effects logit regression models:

$$\text{Model 1: } S_{i,r} = \beta_0 + \beta_1 H_{i,r} + \alpha_i + \epsilon_{i,r}$$

$$\text{Model 2: } S_{i,r} = \beta_0 + \beta_1 H_{i,r} + \beta_2 L_{i,r} + \alpha_i + \epsilon_{i,r}$$

where in Model 1 α_i is a fixed individual effect, $H_{i,r}$ is a vector of indicator variables that take on the value of 1 whenever a particular game history has been achieved and 0 otherwise. Model 2 includes $L_{i,r}$, a vector of indicator variables that interact each game history variable in $H_{i,r}$ with the last ten supergames. The variables take on a value of 1 whenever a particular history has been achieved in the last ten supergames of a supergame type (finite or indefinite) and 0 otherwise. Model 2 is identical to Model 1 in every other respect. We clustered the standard error estimates on sessions so that each session (as opposed to each individual) is treated as a statistically independent observation.

Three types of histories are included in the vector $H_{i,r}$:

(1) The opponent played *D* or *K* in the immediately preceding supergame round: A higher probability of *K* (*D*) in response to a *D* (*K*) in the immediately preceding supergame round would be evidence for play of a "tit-for-tat" strategy. This strategy mimics the action taken by an opponent in the previous round of a game and is

⁸ Each session had fourteen participants, and each session included only five supergames.

possibly empirically relevant (in Axelrod, 1984, this strategy was successful in a tournament).

(2) The opponent played *D* or *K* in any previous round within the supergame: A higher probability of *K* (*D* in response to a *D* (*K*)) in any previous round in a supergame would be evidence for a “grim” strategy, which permanently punishes a defection. This strategy is theoretically relevant, figuring in the construction of cooperative equilibria in repeated games.

(3) Round 3, 4, 5, 6-8, and 9-14 indicators:⁹ In the finitely repeated games conditioning behavior on the round could be empirically relevant, providing evidence that players are learning backward induction strategies (see Selten and Stoeker, 1986, for evidence for this in a prisoner’s dilemma game). In the indefinitely repeated games this type of behavior is not predicted by theory because each supergame continues with a constant and independent probability.

Table 1 presents the regression results. The table is divided into two sections: the left-hand section contains results for the finite games and the right hand section contains results for the indefinite games. Within each section we present coefficient estimates for the game histories, rounds, and interaction terms for both Model 1 (the left-most two columns in the section) and for Model 2 (the right-most two columns in the section). And within each model section we present the results for Player A (the left column) and Player B (the right column).

The regressions show behavior consistent with repeated-game strategies in the finitely repeated games. According to both models, an opponent playing *K* in any previous round decreases the probability that Player A plays *S* (consistent with a “grim trigger” strategy). Reaching rounds four and five also decreases the probability of sending. Play appears to be non-stationary because the round by ten supergames effects are significant and negative: Player A is less likely to play *S* in rounds 4 and 5 in the last ten supergames than in the first ten. Results for Player B are qualitatively similar.

The regressions also show evidence of repeated-game strategies in the indefinitely repeated games. From Model 1 we find evidence that both “grim” and “tit-for-tat” strategies may be relevant strategies for Player A, but we also somewhat surprisingly find evidence for a round effect. Model 2 appears to indicate that this round affect increases over time, as Player A is less likely to play *S* in all rounds (compared with round 2) in the last ten supergames than in the first ten. For Player B, Model 1 appears to imply the same, but Model 2 implies that when we take the interactions with the last half of the supergames into account all behavior appears to be accounted for by the round number, and not by “grim” or “tit-for-tat”. Thus there appears to be an asymmetry in behavior between Player A and Player B in the indefinite games, with Player B apparently conditioning her behaviour increasingly on the round numbers rather than on the history of play with the opponent.

While these results provide evidence for particular types of repeated-game strategies used by the subjects, they are only proxies for drawing such conclusions.

⁹ Round 1 was eliminated due to all lagged independent variables being missing. All other rounds effects are relative to round 2 effects. We grouped round numbers higher than 5 together into two groups because there are relatively few observations of longer games.

Table 1. The effect of game histories on the decision to send and return

Explanatory variables	Finite games		Indefinite games						
	Model 1		Model 2		Model 1		Model 2		
	Player A	Player B	Player A	Player B	Player A	Player B	Player A	Player B	
Opponent's don't send or keep choice									
In the preceding round	0.53	0.16	2.35*	0.31	-0.55**	-0.48*	-0.56*	-0.07	
In any previous round in the supergame	-4.00**	-2.32**	-5.62**	-2.41**	-3.79**	-0.30*	-4.26**	-0.12	
Preceding round by last 10 supergames			-0.11				0.26	-1.07	
Any previous round by last 10 supergames		2.44	note 5				0.42	-0.25	
Round									
Three	-0.58	-1.27	-0.11	-0.93	-0.15	-0.42	0.24	-0.98**	
Four	-1.92**	-3.11**	-1.51*	-3.33**	-0.70**	-1.38*	-0.25	-1.28**	
Five	-2.81**	-4.36**	-2.83**	-4.56**	-0.55	-1.56*	0.41	-1.79**	
Six – eight					-0.65*	-1.32*	0.21	-1.44*	
Nine – fourteen					-1.02**	-1.72*	-0.30**	-2.19**	
Three by last 10 supergames			-0.89**	-0.86**			-0.69**	1.21**	
Four by last 10 supergames			-0.94**	0.51			-0.71**	-0.23	
Five by last 10 supergames			-0.07	0.36			-1.57**	0.44	
Six – eight by last 10 supergames							-1.37**	0.21	
Nine – fourteen by last 10 supergames							-1.06**	1.03	
Constant	2.88**	4.01**	2.72**	4.15**	5.41**	5.32**	5.21**	5.44**	
Number of observations	1820	897	1820	897	2219	1327	2219	1327	
Log-likelihood	-645.13	-332.67	-628.96	-328.39	-708.03	-510.75	-694.68	-502.22	
Log-likelihood test for interaction terms			$p < 0.0001$	$p = 0.1280$			$p = 0.0004$	$p = 0.0170$	

Notes:

1. Significance levels: * indicates 1% level and ** indicates 5% level.
2. Opponent's choice explanatory variables are coded as indicator variables for the following cases.
 - (1) The opponent played Don't Send (Player B) or Keep (Player A) in the previous round.
 - (2) The opponent played Don't Send (Player B) or Keep (Player A) in a previous round of the same supergame.
 - (3) The indicator variable (1) interacted with an indicator for the last 10 supergames.
 - (4) The indicator variable (2) interacted with an indicator for the last 10 supergames.
3. Parameter estimates for fixed subject effects and missing observations are not shown, and Round 1 observations are not included.
4. Standard errors for significance testing are clustered on sessions.
5. This interaction term cannot be estimated due to multicollinearity.

The regressions identify game histories that may figure prominently in decision-making, and they reveal how these histories may affect play on average. In other words, the regressions allow us to infer the type of strategic behavior that occurs on average in the experiments. However, we cannot tell whether behavior is consistent with specific decision rules subjects may have used, or whether they are average responses to potentially many different strategies subjects were using.

Finally, the regressions impose no structure on the type of learning that is occurring, as they simply document whether different types of conditioning variables are becoming more or less relevant to decision making. We cannot, for example, ask the question whether strategic play is moving toward equilibrium, or whether there is a best-response dynamic occurring. To take a step toward answering these types of questions, we now take a look at simple strategies that may be generating these data.

3 The strategy model

Empirical analysis of repeated-game strategies requires identifying an empirically manageable number of strategies to examine from the theoretically large set of repeated-game strategies. Since the set of theoretically possible strategies increases in size exponentially with the number of rounds in repeated games, this set must be reduced for the purpose of strategy inference. The first part of our methodology thus involves a sequence of steps that reduces the theoretically large set. The second part uses data to assess the goodness of fit of the strategies in the reduced set in order to find a best fitting set of strategies.

3.1 *Modelling strategies with finite automata*

There are many contexts in which economists place restrictions on the strategy choices for the solution of complex problems (see Conlisk, 1996, and the references therein). In repeated games, finite automata provide a method for modelling strategies, and in particular they are often used to model the behavior of boundedly rational agents (e.g., Rubinstein, 1986; Abreu and Rubinstein, 1988; Binmore and Samuelson, 1992). Agents are typically assumed to have preferences over both monetary outcomes and the size (i.e., complexity) of their strategy.¹⁰ Finite automata are also useful as inputs to boundedly rational models of learning (e.g., Miller, 1996).

In our case, restricting the strategy set to a set of finite automata provides a useful language with which to describe subject behavior in repeated games; the finite automata enable us to classify and discuss actual sequences of decisions.¹¹ While we do not explicitly test theories of complexity (as in Harrison and McDaniel,

¹⁰ Since complexity of a computational problem is ultimately determined by the time it takes to find a solution, and since the size of an automaton may be an indication of its execution time, size may be a useful measure of its complexity (see Harrison, 2002, for a discussion).

¹¹ Similarly, Spiegler (2005) uses a similar restriction in a theoretical study to be able to specifically discuss testing threats in repeated games.

2002; McKinney and Van Huyck, 2005), inferring finite automata from observed behavior can complement and extend existing theories, thus providing a bridge between repeated-game theory and observed behavior.¹²

While there is substantial *theoretical* research in economics that uses finite automata to model repeated-game strategies, very little empirical evidence exists to support or contradict this research. Empirical evidence for finite automata can therefore complement and extend existing theories; e.g., inferred strategies can help refine equilibrium selection criteria and can also be used with boundedly rational learning models (see Miller, 1986). Inferring finite automata from observed behavior thus helps bridge repeated-game theory with observed behavior.

In this paper we use finite automata to empirically model decision-makers' strategies. Specifically, we use a class of finite automata called Moore Machines (due to Moore, 1956). Each Moore Machine M_i represents a deterministic strategy. Each machine M_i has a finite number of states denoted $q_i^j \in Q_i$. Each state j specifies an action, a_i^j , to take when the machine enters the state, where $a_i^j \in A_i$ is an action available to the machine. The action specified in each state is determined by an output function that maps each state to an action: $\lambda_i : Q_i \rightarrow A_i$. Each state also has a transition function, $\mu_i : A_i \times Q_{-i} \rightarrow Q_i$, which directs the machine to the next state contingent on every possible opponent action. A Moore Machine also designates an initial state, $q_i^1 \in Q_i$, to begin play.

For a few examples of how Moore Machines model strategies, imagine a repeated game with two Players, A and B, and assume each player has two actions in every round of the game. Player A's actions are *S* and *D* and Player B's actions are *R* and *K*. Figure 3 shows twenty-six Moore Machines for Player B. For every machine, each circle represents a state, the notation inside each circle is the state's output function (action) and a double circle indicates the initial state. Arrows represent state transition functions and notation next to arrows indicates the opponent action that triggers the transition.

Strategy M_{b1} represents Player B's strategy to play *R* in every round of the game regardless of Player A's actions. Strategy M_{b3} is a permanent trigger strategy for Player B. It plays *R* in the initial round and continues playing *R* in subsequent rounds as long as Player A plays *S*. If Player A ever plays *D*, however, this action triggers the machine to transition to its second state that specifies playing *K*. Once the strategy enters this second state it remains there for all remaining rounds.

3.2 Constructing the candidate strategy set

Modelling repeated game strategies using Moore Machines is the first step in reducing the set of theoretically possible strategies to an empirically manageable size. To see this, note that Moore Machines make state transitions only in response to the actions of their opponents, but not to their own actions (this type of automaton is

¹² Ken Binmore (Binmore, 1987a,b) suggests that players themselves be modelled as programmed computing machines, thus providing a method to discuss what it means to behave irrationally. Our contribution is to describe the choices of subjects with a set of computational machines with the purpose of better understanding what algorithms may actually be in use in a population of players.

called a “full automaton”; “exact automata” do not have this restriction (e.g., as in Kalai and Sanford, 1988). This restriction seems minimal considering the fact that we will observe repeated games with relatively few repetitions of the stage game, which reduces the number of repeated-game observations in which players do not follow the directions of their own strategy choice.

We perform the second reduction in the strategy set by choosing a candidate set of machines, N , that include strategies that satisfy some initial criteria. For example, the criteria may be based on theory (e.g., equilibria), past empirical evidence (e.g., backward induction or unravelling), a set of machines with specific properties (e.g., a maximum number of states) and/or other priors held by the researcher (e.g., strategies which exhibit behaviors such as reciprocity, biases, heuristics or bounded rationality). In selecting the criteria, the finite automata machine representations are not considered; finite automata are simply tools to represent the behaviors of interest. By choosing to infer a specific candidate set of strategies, we limit the strategies we can infer to this set. However, as we show below, this set can encompass as many (or as few) strategies as the researcher wishes. For instance, in this paper we examine from two to over one thousand strategies in the candidate set using a variety of criteria.

We perform the third reduction of the strategy set by using well known properties of finite automata (e.g., see Hopcroft and Ullman, 1979; Harrison, 1964) to insure that each automaton in the candidate set N is behaviorally unique, i.e., that there exists a sequence of actions that makes possible the empirical identification of each automaton.¹³ The third reduction does not eliminate any strategic behavior that we can infer, it simply eliminates any duplication of identical strategic behavior from the candidate set. For example, a twelve-state machine that indicates a player should play R in each state is behaviorally identical to a one-state machine that indicates playing R , and the third reduction would thus only consider one of these two machines in the candidate set.

The three reductions to form the set of strategies we will empirically consider result in a set of machines, M , that has the following properties. First, M includes exactly one machine representation for each behavior we wanted to examine when we selected the criteria. Second, since each machine in M is behaviorally unique, realizations of actions can occur that allow us to identify each machine from every other one. Third, the identical modeling approach can be used to examine many distinct criteria and many strategies.

Although not to reduce the candidate set, but rather to aid in interpreting the behavior of strategies in the candidate set, we represent each strategic behavior by a machine with the minimum number of states that can represent the behavior. For example, if we wish to include the strategic behavior to unconditionally play action R , then we would include a one-state machine that indicates playing R rather than, say, a two-state or 23-state machine that indicates playing R in every state. By representing every strategic behavior using machines with the minimum number of states, we facilitate interpreting the behavior of the machines.

¹³ The procedure to find the set of behaviorally unique automata involves deriving minimal state representations for each behavior of interest, and then choosing one automaton from each resulting set of isomorphic automata.

4 The estimation procedure

This section describes the estimation procedure to find the best fitting subset of strategies played by a *population* of players from the reduced strategy set M . To avoid over-fitting the data (the goodness of fit of a strategy set weakly increases with the number of strategies in the set), we include a fitness cost that increases with the number of machines. To avoid spurious inference due to serial dependence between supergame observations (i.e., due to multiple supergame observations per subject or possible idiosyncratic learning effects), we test strategy set goodness of fit on a statistically independent hold-out sample of data from a different set of players.

Our procedure will examine *deterministic* strategies, counting the number of supergames the strategies perfectly fit, because the relatively short length of the supergames provides too little information for probabilistic inference. A probabilistic model of strategies would introduce errors in state transitions and action plans; the model would take the form of a markov process with a binomial process generating a pair of actions in each possible state, and could be estimated using hidden Markov, filtering, or Bayesian procedures. We have attempted to estimate such a model but due to the average length of five supergames in our experimental design we have not been able to recover strategies in monte-carlo exercises: the short time series does not provide enough information regarding underlying state transitions.

Our experimental design involved a trade-off between viewing many game histories in which we could assume the underlying strategic behavior is stationary and viewing fewer game histories over a longer time in which play may not be stationary. We chose the former experimental design and were surprised to find that the simple, deterministic strategies fit the vast majority of supergame data.¹⁴ With the current experimental design with many game histories we are also able to recover strategies with monte-carlo exercises that include different types of noise in the data generating process. In fact, the fitness cost we employ in the estimation procedure was directly calibrated from these monte-carlo exercises to account for noisy decision makers.

4.1 Fitting the data

The notion of goodness of fit for finite automata is a simple one; a machine fits a repeated game if, when it replaces the subject, it plays exactly the same actions the subject played. The notion of goodness of fit for a set of machines is similar; a set of machine fits a repeated game if any machine in the set, when it replaces the subject, plays exactly the same actions the subject played. Our goal is to find a subset of machines in M that fits the most data. The problem is that adding a machine to a set weakly increases the set's goodness of fit (we loosely refer to the

¹⁴ With different experimental designs researchers have fit probabilistic models in dynamic contexts: Houser et al. (2004) have subjects play against a stationary strategy played with errors by a computer and use a Bayesian procedure to determine the number and type of decision rules in a dynamic game; Engle-Warnick and Ruffle (2003) extend a procedure first used by El-Gamal and Grether (1995) to a dynamic game to fit probabilistic if-then statements to the actions of buyers in a monopoly game.

number of machines in a set as the set’s complexity). Thus, our goal is to find a subset of machines in M that maximize goodness of fit subject to a cost for the set’s complexity.

To operationalize the inference procedure, we need a few definitions. A repeated game is henceforth referred to as a *supergame*. An *observation* $o_j \in O$ consists of a pair of sequences of actions by opposing players during a supergame. Machine M_i fits supergame o_j if, when M_i replaces the subject in the supergame, M_i responds to the sequence of actions of the subject’s opponent with the exact same sequence of actions taken by the subject it replaced. Let the indicator function $I(o_j, M_i) = 1$ if o_j is fit by M_i and $I(o_j, M_i) = 0$ otherwise.

Definition 1 *The goodness of fit $F(M_i)$ of machine M_i on the set of supergames O is the number of observations $o_j \in O$ that M_i fits:*

$$F(M_i) = \sum_{o_j \in O} I(o_j, M_i).$$

To address heterogeneity that may occur across players (different players may be using different strategies) and over time (individual players may be mixing or learning), we examine the goodness of fit of subsets of n machines, $T_n \subseteq M$. Let $I(o_j, T_n) = 1$ if o_j is fit by at least one machine in T_n and $I(o_j, T_n) = 0$ otherwise.

Definition 2 *The goodness of fit $F(T_n)$ of the set of machines T_n on the set of supergames O is the number of observations $o_j \in O$ that T_n fits:*

$$F(T_n) = \sum_{o_j \in O} I(o_j, T_n).$$

The best fitting set containing n machines, B_n , is found by maximizing the goodness of fit over every possible subset T_n of n machines in M :

$$B_n = \arg \max_{T_n} F(T_n).$$

Note that $F(B_{n+1}) \geq F(B_n)$, since B_{n+1} could always contain all the machines in B_n (note also that B_n is not necessarily a subset of B_{n+1}). Thus the goodness of fit of the best fitting set of machines is weakly increasing with the number of machines in the set. To determine the number of strategies needed in B_n to fit the data and to reduce over fitting the data, we introduce a cost function, $C(n)$, and assume the cost is proportional to the number of strategies in the set. We select the overall best fitting set B from the best fitting sets for each set size, B_n , by maximizing goodness of fit subject to the cost of the number of machines in B_n :

$$B = \arg \max_n f(B_n) - C(n).$$

For simplicity, we let $C(n) = n \cdot c \cdot g(O)$, where $0 < c \leq 1$ and $g(O)$ equals the number of observations in O , so that the marginal cost of including another strategy in B is constant: $C(n + 1) - C(n) = c \cdot g(O)$. To increase the number of strategies in the set by one, the goodness of fit of the best-fitting set must increase

by at least the minimum threshold of $c \cdot g(O)$ (e.g., with $g(O) = 1000$ observations and $c = 5\%$, an additional fifty observations would need to be fit to increase the best fitting set size by one). The appropriate value of c may depend on the number of data generating machines (i.e., degree of heterogeneity), the number of and the actual machines in M and noise in the data generating process. In general, the greater the value of c , the more likely the method will reject actual data generating behavior and the lower the value of c , the more likely the method will spuriously accept non-data generating behavior. We used computer simulations to find a conservative value for c to avoid accepting non-data generating behavior.¹⁵

4.2 Refining the model selection: out of sample goodness of fit

As a second guard against over fitting, we infer the best fitting set in one sample (the *training sample*) and test its goodness of fit on an independent sample (the *test sample*). We use a holdout sample to address the concern that when subjects play many supergames, observations are not independent since there are multiple observations per game and since there is likely to be path dependence in which many players can be affected by common opponents. By requiring the best fitting set of machines to fit an independent sample of data, we mitigate these dependency concerns.

We proceed as follows. If the contribution of any machine in B to the goodness of fit in the test sample is not enough (defined below), we reject the strategy set and move to the best fitting set from the training sample that contains one less strategy. We continue the process until we fail to reject a best fitting set.

The following definition quantifies the contribution of machine M_i to the goodness of fit of set T_n . Let $T_{n,-i}$ denote set T_n excluding M_i , where $M_i \in T_n$.

Definition 3 *The unique goodness of fit $U(T_n, M_i)$ of M_i in the set T_n is the number of supergames o_j that M_i fits and that no other machines in T_n fits:*

$$U(T_n, M_i) = F(T_n) - F(T_{n,-i}).$$

It is easy to show that the unique goodness of fit of each machine in B is at least the threshold level in the training sample; i.e., for all $M_i \in B$, $U(B, M_i) \geq c \cdot g(O)$. We similarly require that the unique goodness of fit of all machines $M_i \in B$ be at least the threshold level in the test sample: $U(B, M_i) \geq c \cdot g(O^T)$, where $g(O^T)$ is the number of observations in the test sample. If the unique goodness of fit of any machine in the test sample falls below this cutoff level, we reject the model $B = B_n$ and select the model $B = B_{n-1}$. We repeat the test until we find a model B^* where all $M_i \in B^*$ pass the threshold criterion for the test sample.

¹⁵ The simulation results are reported in an earlier version of this paper and are available from the authors. The simulation exercises varied the number of data generating machines and included two error specifications.




Figure 2. Player A, set $s \leq 2$

4.3 Selecting the candidate strategy set

We use three criteria to select candidate strategy sets N^A and N^B (superscripts indicate the player type). Our first criterion is to include all strategies in N^A and N^B that have no more than $s = 2$ states. Our motivation for this criterion is primarily to avoid over fitting data, but also to reduce the computational burden and to reflect bounded rationality and complexity.¹⁶ Figures 2 and 3 show sets M^A and M^B that are constructed from sets N^A and N^B using the criterion that no machine has more than $s = 2$ states. There are more machines in M^B than in M^A because whenever Player A plays D Player A's next action is unconditional (since Player A does not observe a Player B action in that round of the game), whereas Player B's action may always be conditional on the action of Player A.

Though the number of machines in M^A and M^B , twelve and twenty-six, respectively, is small, many behaviors are represented. M_{a1} , M_{a2} , M_{b1} and M_{b2} are the unconditional strategies that play S or D for Player A and R or K for Player B. The remaining strategies condition behavior on opponent actions and/or the round. M^A and M^B include the non-cooperative equilibrium pair (M_{a2}, M_{b2}) for the finite and indefinite games and many cooperative equilibrium pairs (e.g., $\{M_{a3}, M_{b1}\}$, $\{M_{a3}, M_{b6}\}$) for the indefinite game. M^A and M^B include trigger strategies such as the extensive form game analogue to the grim-trigger for both players (M_{a3}, M_{b3}) and tit-for-tat for Player B (M_{b5}) . There is no tit-for-tat analogue for Player A in this game.

¹⁶ Bounded rationality suggests a player may not consider all feasible strategies but instead limits himself to "less complex" strategies. The complexity of finite automata machines may be defined in a number of ways (see Osborne and Rubinstein, 1994). One definition is that complexity is positively related to the number of states.

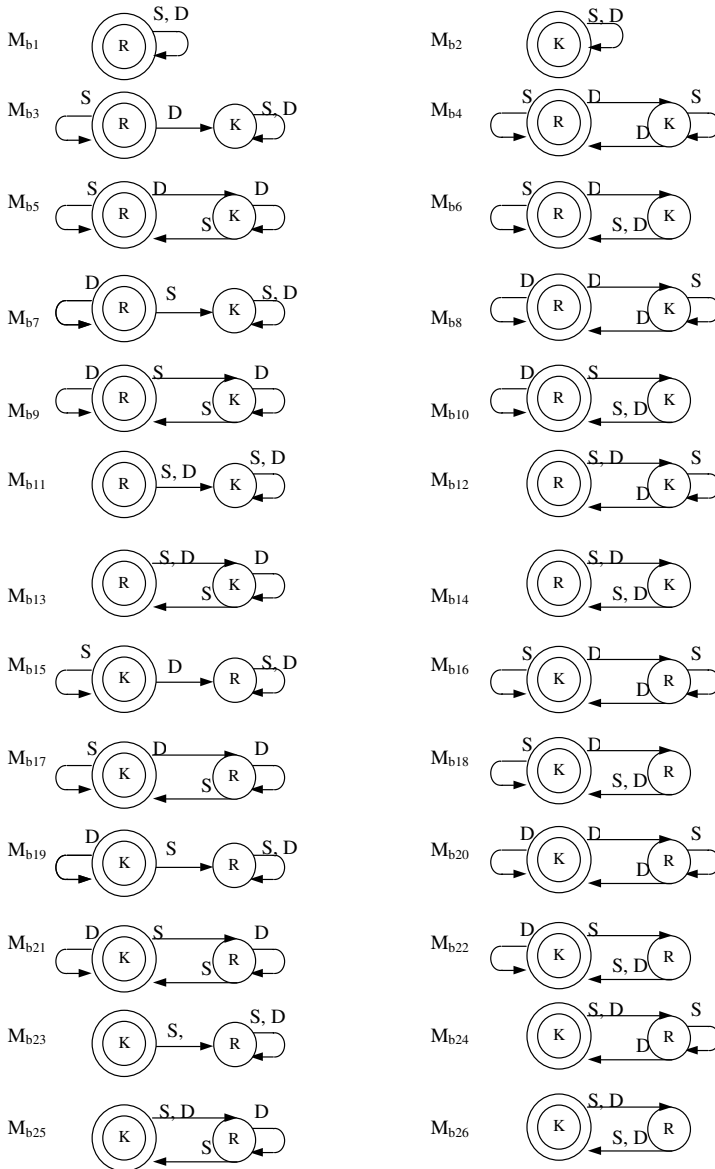




Figure 3. Player B, set $s \leq 2$

Our second criterion is based on theory that suggests different strategies may be used in the finite and indefinite games. In finite games the unique subgame perfect equilibrium is for Player A to never send, M_{a2} , and for Player B to never return, M_{b2} . Past evidence (e.g., Selten and Stoeker, 1986), however, shows that behavior unravels from cooperative behavior towards the non-cooperative equilibrium.¹⁷ To

¹⁷ For example, Player A may play S all five periods and Player B may play R the first four periods, but play K in the fifth period. With experience, Player A may anticipate Player B's behavior and so play



End Game Effect Strategies



Punishment Strategies

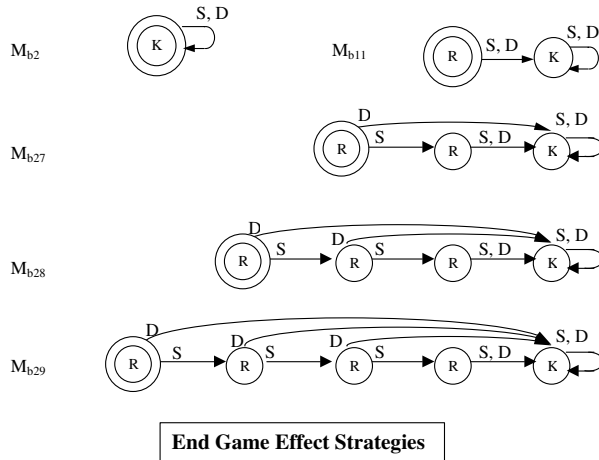
Figure 4. Player A, set “+”

examine this unravelling behavior, we create a set of strategies for each Player A and B that allows unconditional permanent defection from cooperative behavior after each round. Figures 4 and 5 show these strategies for Player A (M_{a2} , M_{a4} , M_{a13} , M_{a14} and M_{a15}) and B (M_{b2} , M_{b11} , M_{b27} , M_{b28} and M_{b29}), respectively.¹⁸

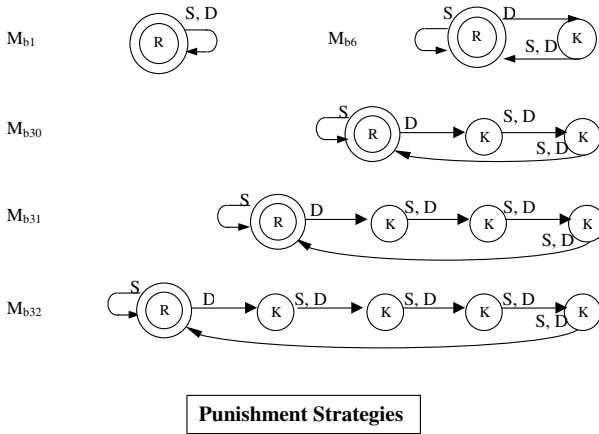
Our third criterion is based on protocols collected in pilot sessions of the indefinite game, which suggest a few additional strategies may be important to fit the data. Subjects indicated they would “punish” an opponent who played his non-cooperative action. The “punishment” involves playing the non-cooperative action

S the first four periods, but then play D in the fifth period. Player B may anticipate this behavior and respond by playing R for only the first three periods and then playing K thereafter. And so on.

¹⁸ For these strategies we assume during the cooperative phase that if an opponent plays his non-cooperative action then the player responds by playing his non-cooperative action thereafter.



End Game Effect Strategies



Punishment Strategies

Figure 5. Player B, set “+”

for a finite number of periods. Figures 4 and 5 show these machines for Player A ($M_{a1}, M_{a5}, M_{a16}, M_{a17}$ and M_{a18}) and B ($M_{b1}, M_{b6}, M_{b30}, M_{b31}$ and M_{b32}), respectively. To analyze the data, we combine the criteria motivated by the finite game unravelling hypothesis and indefinite game protocols to form the “+” sets. We define $s+$ as the union of s (i.e., strategies in M^A and M^B that have no more than s states) and +.

5 Results

Section 5.1 presents criteria for the candidate strategy set that include machines: (1) with a maximum number of states, (2) consistent with past evidence on repeated games and (3) consistent with post-experiment protocol responses. Section 5.2 shows the goodness of fit of these strategy sets and justification for our criteria.

Section 5.3 shows that a small number of inferred machines fit the vast majority of the data and Section 5.4 presents these specific strategies and shows that they are consistent with theoretical predictions and intuition. Section 5.5 shows that strategies inferred over time evolve consistently with best response behavior toward the unique non-cooperative equilibrium in the finite game and the maximally cooperative outcome in the indefinite game.

5.1 Goodness of fit of the candidate sets

The top of each subsection of Table 2 shows the goodness of fit for the candidate sets s and $s+$ for $s = 1, 2,$ and 3 . The four panels show results for Player A and B in the finite and indefinite games. In the finite games the goodness of fit is much higher for set $s+$ than s , holding the maximum number of states constant. For example, sets $1+$ and $2+$ fit over 30% more observations than sets $s = 1$ and $s \leq 2$, respectively, and set $3+$ fits over 17% more observations than set $s \leq 3$. Thus, the $+$ strategies are important to fit finite game data. In indefinite games the $+$ sets have a smaller effect. Although set $1+$ fits almost 20% more observations than set $s = 1$, sets $2+$ and $3+$ fit on average only 5% more observations than sets $s \leq 2$ and $s \leq 3$.

Table 2 provides motivation for examining the $2+$ candidate set. Note that while the number of strategies in M^A increases by a factor of ten from set $2+$ to $3+$ (18 to 180), set $3+$ fits on average only 8% more of the data than set $2+$ in finite or indefinite games. Similarly, while the number of strategies in M^B increases by a factor of more than thirty from set $2+$ to $3+$ (32 to 1058), the $3+$ set fits on average only 4% more of the data than set $2+$. We thus examine the strategies inferred in the $2+$ set; increasing the candidate set to include more strategies (sets 3 or $3+$) does not increase goodness of fit enough to justify the risk of over fitting the data with the ten and thirty fold increases in candidate strategies.

5.2 Number of strategies in the data

The bottom of each subsection of Table 2 shows the goodness of fit of each best fitting set containing $n = 1$ to $n = 6$ strategies for each candidate set. For the most part, our choice of the $2+$ set has little effect on the set goodness of fit, $F(B_n)$; holding n constant, the difference in goodness of fit of the $1+$, $2+$ and $3+$ sets is never greater than 6% (17 out of 280 observations) and is often much less. For example, in finite games for Player B, $F(B_2)$ is 168, 168, and 171 observations for the $1+$, $2+$ and $3+$ sets, respectively. Also note that only a few strategies are necessary to fit a majority of observations. For example, for the $1+$, $2+$ and $3+$ sets $F(B_2) \geq 61\%$ (i.e., the two best fitting strategies fit over 61% of the observations) and $F(B_3) \geq 71\%$.

Using our model selection criterion, Table 2 shows that more heterogeneity (i.e., more strategies) is inferred for Player B than A and more heterogeneity is inferred in the finite than indefinite games. Examining the $2+$ set at a complexity cost of $c = 5\%$, the best fitting sets contain six and three Player B strategies, and three and

one Player A strategies in the finite and indefinite games, respectively.¹⁹ Inferring more Player B than A strategies may be because Player B always observes Player A actions whereas Player A does not always observe Player B actions.

The results in Table 2 are based on the training sample, and we now turn to the test sample to guard against over fitting the data. With one exception we find that in every case (i.e., in the finite and indefinite games for Players A and B) the inferred strategies in the best fitting sets pass the out of sample selection criterion. The exception is in the finite games for Player B; the unique goodness of fit of at least one strategy in the best fitting set in this case is less than $c = 5\%$ out of sample. The best fitting finite Player B model with the most strategies that passes the test sample criterion has $n = 3$ strategies. Therefore, including our out of sample refinement, we infer three Player A and B strategies to fit finite game data and one Player A and three Player B strategies to fit indefinite game data.

5.3 Specific strategies in the data

Figure 6 shows the inferred strategies and Table 3 shows their unique and total goodness of fit in the training and test samples. Henceforth, we refer to each machine's goodness of fit (i.e., the number of machines that it fits, recall Definition 1) as its total goodness of fit in contrast to its unique goodness of fit (i.e., the number of machines that it fits that no other machine in the best fitting set fits, recall Definition 3). In the figures and tables we show the machines in order from the highest to lowest total goodness of fit.

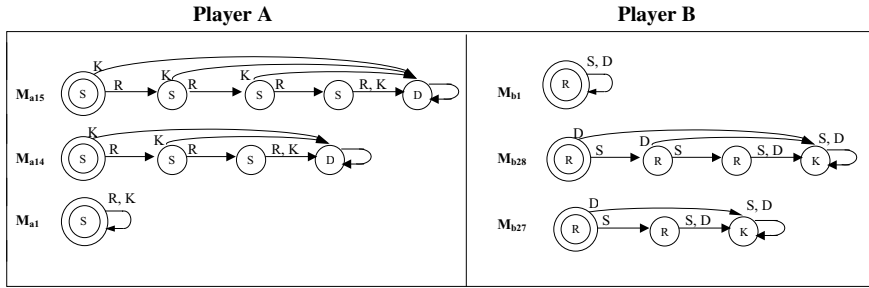
Four of the six inferred finite strategies (M_{a14} , M_{a15} , M_{b27} , M_{b28}) are on the backward induction path. The remaining two finite strategies, M_{a1} and M_{b1} , are the Unconditional Send and Return strategies. The strategies along the backward induction path defect in the fourth or fifth round for Player A and in the third or fourth round for Player B. The multiple inferred strategies along the backward induction path are consistent with at least two hypotheses. First, there may be heterogeneity across players; some players may backward induce one round while others may backward induce two rounds. Second, players may initially backward induce one round and then learn to backward induce two rounds. We examine the second hypothesis below. The key insight is that the method is able to detect heterogeneity.

All inferred finite game strategies are easy to interpret and justify. Four of them are best responses to an opponent's inferred strategy. Player A's Unconditional Send strategy (M_{a1}) is a best response to Player B's Unconditional Return strategy (M_{b1}). Player A's 4th round defection strategy (M_{a14}) is a best response to Player B's 4th round defection strategy (M_{b28}). Player B's 4th round defection strategy (M_{b28}) is a best response to Player A's 5th round defection strategy (M_{a15}) and Player B's 3rd round defection strategy (M_{b27}) is a best response to Player A's 4th round defection strategy (M_{a14}).


¹⁹ Computer simulations replicating game conditions indicate that $c = 5\%$ is conservative in the sense that over fitting the data did not occur. The simulations also validated the ability to recover the known data generating machines in the presence of noise.

Table 2. Training sample goodness of fit of base and best-fitting machines

Finite games													
A Players					B Players								
Goodness of fit of base sets					Goodness of fit of base sets								
Set	No. of machines	Goodness of fit	% Fit		Set	No. of machines	Goodness of fit	% Fit					
1 state	2	70	25.0%		1 state	2	142	51.6%					
2 state	12	171	61.1%		2 state	26	174	63.3%					
3 state	176	206	73.6%		3 state	1054	208	75.5%					
1 state +	10	255	91.1%		1 state +	10	265	86.4%					
2 state +	18	256	91.4%		2 state +	32	268	97.5%					
3 state +	180	278	99.3%		3 state +	1058	274	99.6%					
Goodness of fit of best-fitting sets B_n					Goodness of fit of best-fitting sets B_n								
No. of machines	Base set for B_n						No. of machines	Base set for B_n					
	1	2	3	1+	2+	3+		1	2	3	1+	2+	3+
1	65	145	145	145	145	145	1	121	124	126	124	124	126
2	70	156	169	197	197	205	2	142	153	157	168	168	171
3	n.a.	162	182	227	227	239	3	n.a.	172	183	196	196	199
4	n.a.	167	189	239	239	253	4	n.a.	174	203	222	223	225
5	n.a.	170	194	245	245	259	5	n.a.	174	205	246	247	249
6	n.a.	171	197	250	250	264	6	n.a.	174	206	265	266	269
280 total observations					275 total observations								
Indefinite games													
A Players					B Players								
Fitness of base sets					Fitness of base sets								
Set	No. of machines	Goodness of fit	% Fit		Set	No. of machines	Goodness of fit	% Fit					
1 state	2	170	60.7%		1 state	2	187	66.8%					
2 state	12	227	81.1%		2 state	26	225	80.4%					
3 state	176	249	88.9%		3 state	1054	244	87.1%					
1 state +	10	224	80.0%		1 state +	10	240	85.7%					
2 state +	18	237	84.6%		2 state +	32	251	89.6%					
3 state +	180	255	91.1%		3 state +	1058	261	93.2%					
Goodness of fit of best-fitting sets B_n					Goodness of fit of best-fitting sets B_n								
No. of machines	Base set for B_n						No. of machines	Base set for B_n					
	1	2	3	1+	2+	3+		1	2	3	1+	2+	3+
1	170	202	202	193	202	202	1	158	161	162	161	161	162
2	170	211	212	206	211	212	2	187	192	194	188	192	194
3	n.a.	218	222	213	218	222	3	n.a.	212	218	208	212	218
4	n.a.	222	227	216	222	227	4	n.a.	217	230	221	224	231
5	n.a.	224	231	218	225	231	5	n.a.	221	237	230	233	243
6	n.a.	226	235	220	227	237	6	n.a.	222	240	236	239	249
280 total observations					280 total observations								



a



b

Figure 6. **a** Inferred strategies in all finite supergames. **b** Inferred strategies in all indefinite supergames

The remaining two inferred finite game strategies are also easily justified. Although Player A’s 5th round defection strategy (M_{a15}) is not a best response to any of Player B’s inferred strategies, it may reflect Player A anticipating Player B playing K in the last round; Player B on average plays K 44% of the time in the last round.²⁰ And although Player B’s Unconditional Return strategy (M_{b1}) is also not a best response to any of Player A’s inferred strategies, it reflects a preference for fairness or equity consistent with utility functions proposed by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000).

In the indefinite game the only Player A strategy inferred, M_{a3} , is the analogue to the Grim Trigger strategy in the repeated Prisoner’s Dilemma game. Strategy M_{a3} fits 72% and 67% of the observations in the training and test samples, respectively. This strategy is a best response to Player B’s inferred strategy M_{b6} . In fact, M_{a3} and M_{b6} form a cooperative equilibrium pair. The two other inferred Player B strategies, M_{b11} and M_{b25} , are not best responses to M_{a3} . However, given that 28% of Player A observations are not consistent with the Grim Trigger (M_{a3}), M_{b11} and M_{b25} may be justified in terms of profit maximization.

Figure 7 shows the total goodness of fit (combining the training and test samples), for the first and last ten supergames separately, for inferred strategies that are best responses to at least one inferred opponent strategy. In the finite game, the total goodness of fit of the two inferred Player A best response strategies (M_{A1} and

²⁰ All best fitting strategy sets for Player B with more than three strategies include the 5th round defection strategy M_{b29} .

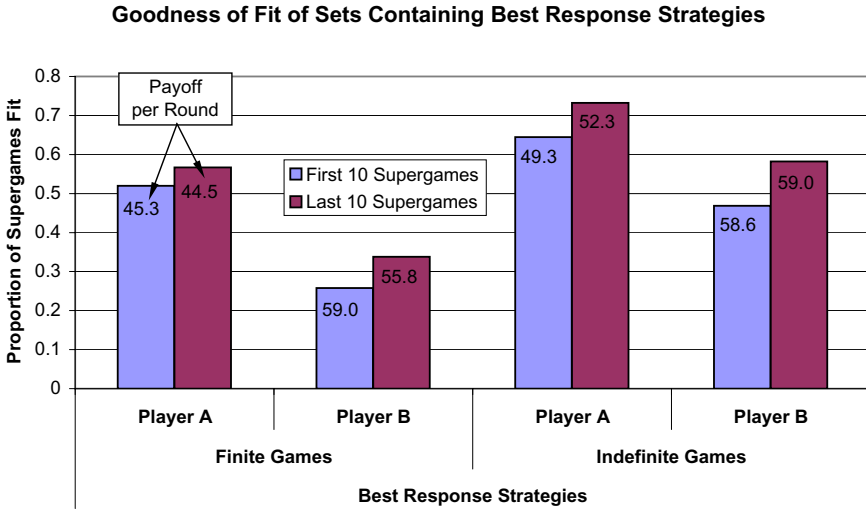


Figure 7. Goodness of fit and average payoff of inferred best response strategies

M_{A14}) increases by 5% (from 52% to 57%), and the total goodness of fit of the inferred Player B best response strategy (M_{B28}) increases by 8%. In the indefinite game, the total goodness of fit of the inferred best response Player A and Player B strategies increases by 9% and 11%, respectively. The increased goodness of fit of the inferred strategies over time suggests that players are increasingly best responding to the strategic behavior of their opponents.

Figure 7 also presents the average per round payoff received by playing the inferred best response strategies.²¹ In the finite game, the average per round payoff decreases for both players, moving toward the unique non-cooperative equilibrium per round payoff of \$0.40. For instance, Player B’s average per round payoff decreases from \$0.590 to \$0.558. In the indefinite game, the average per round payoff increases for both players, moving toward to the cooperative equilibrium per round payoff of \$0.60. For instance, Player A’s average per round payoff increases from \$0.493 to \$0.523. Thus, not only are the inferred best response strategies fitting an increasing proportion of supergames over time, but payoffs are also moving towards equilibrium payoffs in both finite and indefinite games.

In sum, inferred strategies differ across conditions in behaviorally meaningful ways. Four of the six inferred finite game strategies reflect end game effects on the backward induction path while none of the inferred indefinite game strategies reflect end game effects. Further, four of the six finite game strategies are best responses to inferred opponent strategies and one more reflects preferences modeled by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000)). Two of the inferred strategies in the indefinite game form a cooperative equilibrium pair. Finally, in both the finite and indefinite games, the total goodness of fit of inferred strategies that are a best response to inferred opponent strategies increased over time, and the average

²¹ This payoff is calculated by determining the average payoff per round for each supergame the strategy fits, and then averaging across all supergames fit by the strategy.

Table 3. Inferred strategies and their goodness of fit

		Machine	Training sample		Test sample	
			Unique fit	Total fit	Unique fit	Total fit
Finite games	Player A	M_{a15}	20%	40%	15%	42%
		M_{a14}	17%	38%	8%	35%
		M_{a1}	23%	23%	9%	9%
	Player B	M_{b1}	20%	44%	12%	29%
		M_{b28}	15%	41%	13%	43%
		M_{b27}	10%	19%	17%	38%
Player A	M_{a3}	72%	72%	67%	67%	
Indefinite games	Player B	M_{b6}	43%	57%	32%	46%
		M_{b11}	7%	22%	5%	21%
		M_{b25}	11%	11%	15%	17%

Notes:

1. Bold face indicates strategies that are best responses to inferred opponent strategies.
2. Calculations based on $c = 5\%$ threshold.

payoff received from playing these strategies moved toward the average payoff for an equilibrium payoff for each game.

5.4 Behavior over time

In this section we examine whether inferring multiple strategies is due to players adapting different strategies over time. We address this question by inferring strategies for the first and last ten supergames. Figures 8 and 9 show the inferred finite and indefinite game strategies, respectively, for the first and last ten supergames and Table 3 shows the goodness of fit of these strategies.

In the finite game, the same strategy sets that are inferred for all twenty supergames for Player A and Player B are inferred for the first ten supergames (see Fig. 8a for Player A and Fig. 8b for Player B). However, for Player A for the last ten supergames the five-round counter (M_{a15}), which was not a best response to any inferred opponent strategy, and the always send (M_{a1}) machines are no longer inferred. On the other hand, the four round counter (M_{a14}), which is a best response to the Player B four round counter (M_{b28}), fits increasingly more of the data over time. This change in behavior likely reflects Player A backward inducing to a greater degree. For Player B, all three inferred strategies for the first ten supergames are replaced with two new strategies for the last ten supergames. First, Player B may be replacing the always return strategy (M_{b1}) with a similar strategy (M_{b30}) that returns so long as Player A plays S but plays K for two rounds if Player A plays D , and then starts over again (subjects in the pilot sessions articulated a similar strategy). In other words, Player B appears to be substituting unconditional reciprocity for a two round trigger punishment strategy. Second, Player B may be

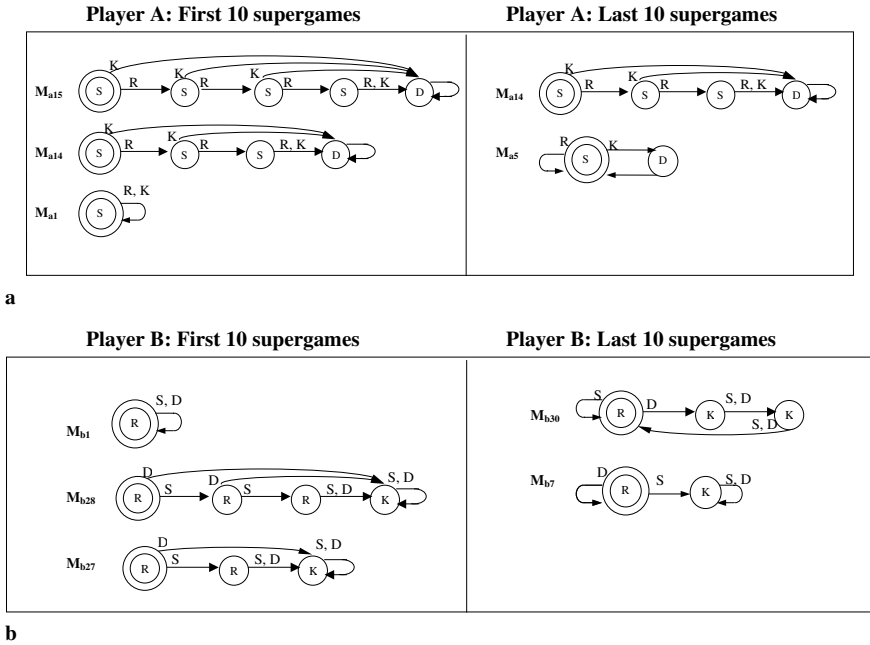


Figure 8. **a** Inferred player A strategies in finite supergames. **b** Inferred player B strategies in finite supergames

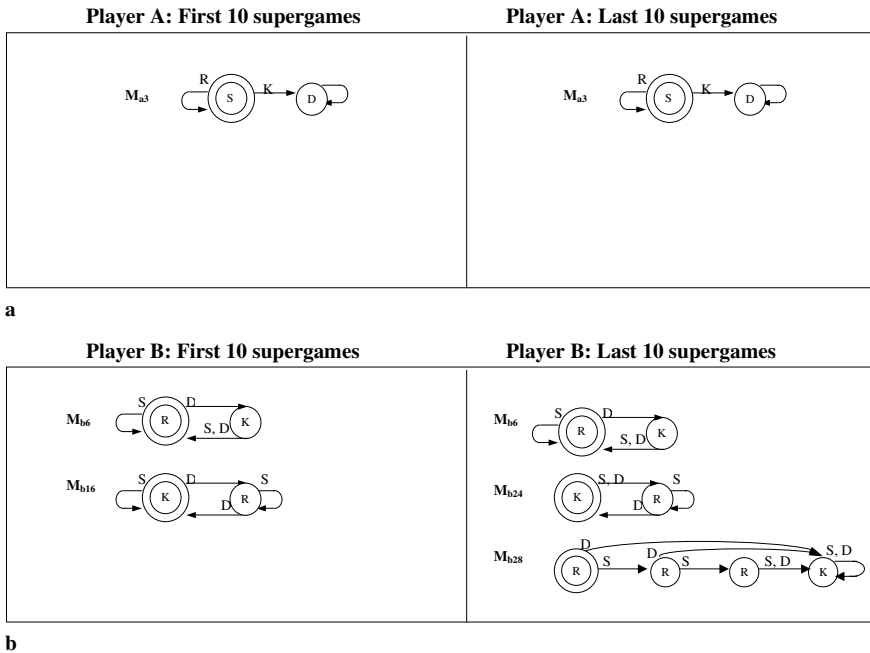


Figure 9. **a** Inferred player A strategies in indefinite supergames. **b** Inferred player B strategies in indefinite supergames

Table 4. Inferred strategies and their goodness of fit

		First ten supergames						Last ten supergames											
		Machine			Training sample			Test sample			Machine			Training sample			Test sample		
		Unique fit	Total fit	Total fit	Unique fit	Total fit	Total fit	Unique fit	Total fit	Total fit	Unique fit	Total fit	Total fit	Unique fit	Total fit	Total fit	Unique fit	Total fit	Total fit
Finite games	Player A	M _{a,15}	28%	42%	21%	47%	M _{a,15}	46%	46%	46%	46%	34%	34%	34%	34%	34%	34%	34%	34%
		M _{a,14}	15%	29%	10%	36%	M _{a,5}	27%	27%	27%	27%	6%	6%	6%	6%	6%	6%	6%	
		M _{a,1}	26%	26%	13%	13%	→	→	→	→	→	→	→	→	→	→	→	→	→
	Player B	M _{b,1}	24%	43%	15%	31%	M _{b,30}	41%	46%	46%	46%	25%	34%	34%	34%	34%	34%	34%	34%
		M _{b,28}	20%	42%	18%	44%	M _{b,7}	17%	21%	21%	21%	7%	15%	15%	15%	15%	15%	15%	15%
		M _{b,27}	11%	17%	15%	30%	↗	↗	↗	↗	↗	↗	↗	↗	↗	↗	↗	↗	↗
Indefinite games	Player A	M _{a,3}	67%	67%	62%	62%	M _{a,3}	77%	77%	77%	77%	71%	71%	71%	71%	71%	71%	71%	
	Player B	M _{b,6}	54%	54%	62%	62%	M _{b,6}	28%	61%	61%	61%	17%	55%	55%	55%	55%	55%		
		M _{b,16}	14%	14%	20%	20%	M _{b,24}	9%	9%	9%	9%	13%	15%	15%	15%	15%	15%		
						M _{b,28}	9%	42%	42%	42%	11%	51%	51%	51%	51%	51%	51%		

Notes:

1. Calculations based on $c = 5\%$ threshold.
2. Arrows indicate machines with nearly identical structures, differing at most by two transitions.

replacing the three and four round conditional counters (M_{b27} and M_{b28}) with a strategy (M_{b7}) that is functionally similar to a two period counter.

In the indefinite games (see Fig. 9a for Player A and Fig. 9b for Player B) we continue to only infer the Grim Trigger strategy for Player A (M_{a3}). Averaging across the training and test samples, note that over time Player A increasingly makes choices consistent with this strategy; its goodness of fit rises from 65% in the first ten to 74% in the last ten supergames. Likewise, Player B increasingly makes choices consistent with machine M_{b6} , fitting 44% in the first 10 and 58% in the last ten supergames. Thus, the inferred equilibria pair (M_{a3} , M_{b6}) over all twenty supergames is inferred in both the first and last ten supergames and, as already noted, these strategies increasingly fit the data over time. In the last ten indefinite games the four round counter (M_{b28}) is also inferred and fits over 40% of the supergames. This behavior may reflect a form of gambler’s fallacy in which Player B incorrectly anticipates that the supergame relationship is increasingly likely to end after the fourth round. The other Player B strategies we infer in the first and last ten supergames, M_{b16} and M_{b24} , respectively, play K in the first round. Note that we also infer a Player B strategy that plays K in the first round using all twenty supergames (M_{b25}). Thus, it appears that to describe Player B’s behavior, we must include a strategy that plays K in the first round, but it is not clear what behavior occurs once he plays K .²²

In sum, behavior in the finite and indefinite games evolves in a best-response manner. In the indefinite game, both players increasingly play strategies consistent with a cooperative equilibrium. In the finite game Player A increasingly backward induces and Player B adapts a two-round trigger punishment strategy. Thus, the heterogeneity observed over all twenty supergames may be partially explained by players adapting different strategies over time.

6 Validity of the procedure

We contribute to our knowledge of play in repeated games with a model of strategies, an inference procedure, and an experimental design. We first presented evidence with a fixed effects panel model that game histories affect play in repeated games. We developed an inference procedure that uncovered simple models of computation as likely strategies behind the actions of subjects in the experiments.

The difficulty with inferring strategies from actions is that strategies are unobserved, and many different strategies could have generated any sequence of observed actions. To overcome this difficulty we made modelling decisions that traded off our ability to use conventional means of inference with our ability to observe play in many repeated games. Specifically, we chose an average length of five rounds per supergame so that we could observe many game histories to help

²² We may infer different behavior for Player B after he plays K across the first and last ten supergames as well as across all twenty supergames because we rarely observe this path of play. To see this, recall Player A’s Grim Trigger strategy (M_{A3}) fits 72% of all supergames, so in only 28% of the supergames in which Player B plays K do we observe him taking another action. And since Player B only plays K in the first round in 11% of the supergames, we only observe Player B’s behavior after he plays K in 3% of the supergames.

us uncover strategies. And since these relatively short supergames provided relatively little information regarding the transition of play from one state to another we employed deterministic finite automata to explain our data.

We think it is extraordinary that such simple, deterministic models (1) explained the vast majority of supergame play, (2) dramatically confirmed theoretical predictions in repeated games and (3) provided such a clear and coherent characterization of play. Not only do inferred strategies seem to be evolving toward equilibrium strategies in both cases, the heterogeneity that exists provides evidence that subjects are best responding to play in our dynamic framework. With these results in hand, we can explore longer repeated games, introducing errors into the finite automata, collecting more state transition information to do formal inference. We can do so with the confidence that relatively few basic strategies seem to be behind the actions of subjects in repeated games. The inferred strategies are consistent with the panel data results and provide additional insights to our understanding of repeated-game behavior of the subjects. Subjects are mostly making choices consistent with just one to three specific repeated game strategies and the choice of strategies evolved over time in an apparently best response manner.

7 Conclusion

This paper is a step toward bridging the gap between theory and empirical observation in repeated games. Repeated-game strategy inference will allow researchers to form and examine hypotheses regarding strategies based on empirical observation. In our case, evidence for the use of relatively few strategies that are best responses to opponent strategies contributes to the refinement of equilibria in the repeated game.

In this paper we combine a strategy model, an inference procedure and an experimental design to infer unobserved repeated-game strategies from observed actions. To demonstrate the inference method, we examine finitely and indefinitely repeated trust games. In finite games we find evidence of players using strategies with end-game effects. In indefinite games we find substantial evidence for a harsh trigger strategy for Player A. The punishment phase of this trigger strategy is harsh enough that the strategy may be included in the construction of repeated-game equilibria. Further, in both finite and indefinite games only a small number of strategies are needed to fit the vast majority of the data, most of the inferred Player A and B strategies are best responses to one of the inferred strategies of their opponents and the inferred strategies evolve in a best-response manner. And the data we do not explain could represent either trembles or rarely-used strategies.

More generally, by starting to bridge the gap between observed behavior and theoretical predictions of play in repeated games, this paper moves researchers closer to applying similar models to behavior in the field. The next step will require introducing errors in the strategies for the purpose of fitting them to more general decision-making environments than those found in the lab. As such, this paper is a necessary foundation in a larger research agenda; if we fail to find evidence for the repeated game strategy model in our experimental environment, then the application of the strategy model to less controlled environments could be called into question.

However, since a small number of deterministic strategies fit the vast majority of decision-making events (repeated games) in our experiment, we have confidence to move forward. In related work making advances in this direction, Engle-Warnick and Ruffle (2003) fit probabilistic if-then statements to decision making in market experiments. We are also in the process of applying filtering techniques from time series analysis to the strategy inference problem, taking observed actions as noisy representations of unobservable states, in a direct extension of the finite automata model in this paper.

Although we think introducing errors into the empirical analysis to examine decision-making rules is an important next step in this research agenda, we are cautious to note that introducing errors will likely change the behavioral interpretation of the strategies in fundamental ways. For example, two harsh trigger strategies that exhibit a non-zero probability of making errors when transitioning ultimately end up jointly in their punishment states with high probability, a result that is in direct contrast with the deterministic case. When decisions are made for longer periods of time, more complex models (e.g., learning models and models that allow switching strategies over time) will be necessary. Our study provides a starting point for choosing which strategies to include in these more complex models. For these reasons, the focus in this paper on deterministic strategies and the finding that they describe much of the data, is an important foundation for future research.

References

- Abreu, D., Rubinstein, A.: The structure of Nash equilibrium in repeated games with finite automata. *Econometrica* **56**, 1259–1282 (1988)
- Axelrod, R.: The evolution of cooperation. New York: Harper Collins 1984
- Berg, J., Dickhaut, J., McCabe, K.: Trust, reciprocity, and social history. *Games and Economic Behavior* **10**, 278–305 (1995)
- Binmore, K.: Modeling rational players I. *Economics and Philosophy* **3**, 9–55 (1987)
- Binmore, K.: Modeling rational players II. *Economics and Philosophy* **4**, 179–214 (1987)
- Binmore, K., Samuelson, L.: Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory* **57**, 278–305 (1992)
- Bolton, G., Ockenfels, A.: ERC: a theory of equity, reciprocity, and competition. *The American Economic Review* **90**, 166–193 (2000)
- Brandts, J., Charness, G.: Hot vs. cold: sequential responses and preference stability in experimental games. *Experimental Economics* **2**, 227–238 (2000)
- Conlisk, J.: Why bounded rationality? *Journal of Economic Literature* **34**, 669–700 (1996)
- Costa-Gomes, M., Crawford, V., Broseta, B.: Cognition and behavior in normal-form games: an experimental study. *Econometrica* **69**, 1193–1237 (2001)
- El-Gamal, M., Grether, D.: Are people Bayesian? Uncovering behavioral strategies. *Journal of the American Statistical Association* **90**, 1137–1145 (1995)
- Engle-Warnick, J.: Inferring strategies from actions: a nonparametric, binary tree classification approach. *Journal of Economic Dynamics and Control* **27**, 2151–2170 (2003)
- Engle-Warnick, J., Ruffle, B.: The strategies behind their actions: a method to infer repeated-game strategies and an application to buyer behavior. Nuffield College Working Paper 2001-W28 (2001)
- Fehr, E., Schmidt, K.: A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics* **114**, 817–868 (1999)
- Fudenberg, D., Maskin, E.: The Folk theorem of repeated games with discounting or with incomplete information. *Econometrica* **125**, 533–554 (1986)
- Harrison, G.: Experimental behavior as algorithmic process: an introduction. Working Paper, University of Central Florida (2002)

- Harrison, G., McDaniel, T.: Voting games and computational complexity. Working Paper, University of Central Florida (2002)
- Harrison, G.: Introduction to switching and finite automata. Reading: Addison-Weseley 1964
- Hopcroft, J., Ullman, J.: Introduction to automata theory, languages, and computation. Reading: Addison-Weseley 1979
- Houser, D., Keane, M., McCabe, K.: Behavior in a dynamic decision problem: an analysis of experimental evidence using a Bayesian type classification algorithm. *Econometrica* **72**, 781–822 (2004)
- Johnson, E., Camerer, C., Sen, S., Rymon, T.: Detecting failures of backward induction: monitoring information search in sequential bargaining. *Journal of Economic Theory* **104**, 14–47 (2002)
- Kalai, E., Stanford, W.: Finite rationality and interpersonal complexity in repeated games. *Econometrica* **56**, 397–410 (1988)
- Kreps, D.: Corporate culture and economic theory. In: Alt, J., Shepsle, K. (eds.) *Perspectives on positive political economy*, pp. 90–43. Cambridge: Cambridge University Press 1990
- McKinney, N., Van Huyck, J.: Estimating bounded rationality and pricing performance uncertainty. *Journal of Economic Behavior and Organization* (2005) (forthcoming)
- Miller, J.: The coevolution of automata in the repeated prisoner's dilemma. *Journal of Economic Behavior and Organization* **56**, 87–112 (1996)
- Moore, E.: Gedanken experiments on sequential machines. In: *Automata studies*. Princeton: Princeton University Press 1956
- Osborne, M., Rubinstein, A.: *A course in game theory*. Cambridge: MIT Press 1994
- Rubinstein, A.: Finite automata play the repeated prisoner's dilemma. *Journal of Economic Theory* **39**, 83–96 (1986)
- Selten, R., Mitzkewitz, G., Urlich, R.: Duopoly strategies programmed by experienced traders. *Econometrica* **65**, 517–555 (1997)
- Selten, R., Stoeker, R.: End behavior in sequences of finite prisoner's dilemma supergames. *Journal of Economic Behavior and Organization* **65**, 47–70 (1986)
- Slonim, R.: Learning and bounded rationality in a search-for-the-best alternative. *Journal of Economic Behavior and Organization* **25**, 141–165 (1994)
- Speigler, R.: Testing threats in a repeated game. *Journal of Economic Theory* **121**, 214–235 (2005)
- Stahl, D., Wilson, P.: On players' models of other players: theory and experimental evidence. *Games and Economic Behavior* **10**, 218–254 (1995)
- Van Huyck, J., Battalio, R., Walters, M.: Is reputation a substitute for commitment in the peasant-dictator game? Unpublished manuscript, Economics Department, Texas A & M University (2001)