

# Inferring Repeated Game Strategies From Actions: Evidence From Trust Game Experiments\*

Jim Engle-Warnick  
Nuffield College  
Oxford University  
Oxford OX1 1NF, UK  
jim.enge-warnick@nuff.ox.ac.uk

Robert L. Slonim<sup>†</sup>  
Department of Economics  
Weatherhead School of Management  
Case Western Reserve University  
Cleveland, OH, 44106  
rls18po.cwru.edu

July 20, 2001

## Abstract

This paper is an empirical study, using new experimental data, of repeated game strategies in trust games; its goal is to identify strategies that people use in repeated games. We develop a strategy inference method that maps observed actions to a set of best fitting unobserved repeated game strategies. Data analysis shows the ability of the method to infer distinct but intuitive and theoretically justified sets of strategies across finitely and indefinitely repeated games. In indefinitely repeated trust games we infer trigger strategies that are consistent with equilibria. In finitely repeated games we infer strategies with end-game effects. Almost all strategies inferred are best responses to the inferred strategies of opponents. For the first time we hypothesize repeated game strategies based on observed behavior, and characterize observed behavior using the core game theory concept of repeated-game strategies.

*JEL classification: C72, C80, C90*

*Keywords: Game Theory, Empirical Methods, Experimental Economics, Repeated Games, Trust*

---

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

<sup>†</sup>Corresponding Author: Tel: (216) 368-5811, Fax: (216) 368-5039.

# 1 Introduction

Social scientists use repeated games as metaphors to study many environments such as investment decisions, buyer-supplier and employee-employer relationships, public goods and common pool resource problems, the nuclear arms race, trench warfare and voting (e.g., Ostrom (1998) and Axelrod (1984)). Although theoretical analysis of repeated games is extensive, empirical analysis of the strategies people use in them has been limited due to a lack of data and a lack of a strategy inference method. Experimental techniques have developed to provide repeated game data (e.g., Selten and Stoeker (1986) and Van Huyck, Battalio, and Walters (1997)). This paper presents a method to infer strategies that people use in repeated games. The method allows researchers to develop an empirically based strategy model that grounds the core game theoretic concept of a strategy in observed behavior. The method is a step toward better understanding decision-making in repeated games. A better understanding can result, for example, from being able to empirically identify a subset of equilibria in an environment where there are many.

Our empirical method maps observed actions in a repeated game to unobserved strategies. The need for a method to infer strategies from actions exists because strategy choices are not observed, and cannot be directly recovered, in most field settings. For example, researchers can observe Federal Reserve actions (raise, maintain or lower the discount rate), union actions (go on strike or don't go on strike) or a firm's actions (increase or decrease orders), but cannot observe the underlying strategy choices that generate the actions.

Inferring strategies from actions involves overcoming the problem that any observed history of actions can be the outcome of many strategies. For example, if we wish to understand someone's plan (i.e., strategy) as to when to carry an umbrella, we would have had a difficult time inferring the plan if we observed that the person always carried an umbrella in Oxford in the fall of 2001 (when it rained constantly). Is the plan to always carry an umbrella regardless of the weather? Is the plan to only carry one whenever it rains in the morning? Is the plan to carry one whenever

it is cloudy? We reduce the severity of this problem by running experiments to observe behavior in many repeated games, thus inducing potentially many game histories. Continuing the example, by examining potentially many histories, we hope to observe actions on rainy, cloudy and sunny days.

This paper uses the inference method to study strategies used in repeated play of a trust game. In the 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.

Using the strategy inference method, we find that subjects play trigger strategies in indefinite games that are consistent with equilibria, and play strategies with end-game effects in finite games. In the finite game we also find that the strategies we infer evolve over time in a manner consistent with unraveling behavior. Almost all inferred strategies can be justified as best responses to the strategies of the opponents. Thus, the analysis shows the robustness of the methodology to find intuitive and theoretically justified results.

Strategy inference complements existing approaches to study repeated game strategies. One approach, called the strategy method, creates an environment where the researcher can directly observe subjects' strategy choices. In this environment, players choose strategies that are then played against each other by the researcher (Selten and Mitzkewitz (1997)).<sup>1</sup> In our method, strategy choices are unobserved so that the decision process subjects face more closely reflects decision-making in field settings. Another approach is to assume *a priori* that people use only a few strategies or a specific class of strategies (e.g., Selten and Stoeker (1986)). Our method imposes fewer *a priori* assumptions regarding the number and types of strategies

---

<sup>1</sup> Slonim (1994) and Brandts and Charness (2000) compare the strategy and direct method.

that researchers can investigate.

Strategy inference can be thought of as part of a more general line of research that seeks to understand decision rules that people use in games (e.g., Costa-Gomez, Crawford, and Broseta (2001), El-Gamal and Grether (1995), Engle-Warnick (2001), Slonim (1999), Stahl (2000)). Similar to this research, we infer unobserved strategies from observed actions, but unlike it our context is repeated games.

The remainder of the paper describes the strategy model, the method of inference and the application to the trust game experiment. Section 2 presents the strategy model. Section 3 shows how we infer strategies from observed actions. Section 4 describes a new trust game experiment. The experiment shows the robustness of the method to infer strategies in both indefinite and finite games and across two player types. Section 5 presents the logical and distinct strategies inferred. The conclusion offers directions for future research and extensions to field research.

## 2 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 fitness of the strategies in the reduced set in order to find a best fitting set of strategies.

### 2.1 Modelling Strategies with Finite Automata

We model strategies with finite automata. Finite automata are one of several ways to model repeated game strategies. A finite automaton is an algorithm that receives a sequence of actions taken by an opponent in a repeated game, one round at a time, and produces an action to play each round. We use finite automata to model

repeated game strategies for several reasons: (1) they are well suited to recognize patterns in data, (2) they can represent interpretable behavior (e.g., reciprocity), (3) their theoretical properties are well understood in economics (e.g., Abreu and Rubinstein (1988) and Binmore and Samuelson (1992)) and (4) they are useful as inputs to computational economic models (Miller (1996)).

We use a class of finite automata called Moore Machines (Moore (1956)). Each Moore Machine  $M_i$  represents a deterministic strategy. Each machine  $i$  has a finite number of states denoted  $q_i^j \in Q_i$ . Each state  $j$  specifies an action,  $a_i^j$ , to take (e.g., carry an umbrella, don't carry one) 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, suppose there is 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 2 shows 26 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.

Selecting Moore Machines to model repeated game strategies reduces the set of theoretically possible strategies. To see this, note that Moore Machines do not

provide a contingent plan of action in the event that they themselves make a mistake, i.e., they make their state transitions based on opponent actions but not on their own actions; thus they are a subset of the repeated game strategy space. Thus, selecting Moore Machines to represent repeated game strategies is the first step we take to reduce the set of theoretically possible strategies to an empirically manageable size.

## 2.2 Constructing the Candidate Strategy Set

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 unraveling), 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. However, it is possible that the criteria selected may lead to multiple machine representations that produce the identical sequence of actions for every possible sequence of opponent actions. For example, suppose one criterion is to include strategic behavior that is not very complex, and to implement this criterion we define strategic complexity as those strategies that can be represented by machines with two or fewer states.<sup>2</sup> This set would include a machine with two states that takes the same action in each state and the set would also include a machine with one state that takes the same action. So, no matter what sequence of actions an opponent takes, these two machines will always produce the same sequence of actions.

If multiple machines produce an identical sequence of actions for every sequence of opponent actions, (henceforth referred to as behaviorally equivalent machines), then no sequence of opponent actions can occur that would let us empirically dis-

---

<sup>2</sup> See Kalai and Stanford (1988) for an example of the use of minimal state representations in a model of complexity.

tinguish between the machines. We thus perform the third reduction by eliminating behaviorally equivalent machines. In particular, if multiple machines are behaviorally equivalent, we keep the machine with the fewest states.<sup>3</sup>

The modelling process results 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 no machines in  $M$  are behaviorally equivalent, realizations of actions can occur that allow us to identify each machine from every other one. Third, since every behavior is represented using a machine with a minimum number of states, behaviorally interpreting machines is facilitated. Finally, the identical modeling approach can examine many criteria. To demonstrate this flexibility, we examine several criteria that produce sets including from 2 to over 1000 strategies.

### 3 The Estimation Procedure

This section describes our estimation procedure to find the best fitting subset of strategies within the reduced set  $M$ . Since the fitness of a subset of automata weakly increases with the number of machines in the set, we take great care not to over fit the data. To avoid over fitting, we take two steps: (1) we include a fitness cost that increases with the number of machines in a set, and (2) after finding the best fitting subset of machines, we test this fitness on a hold-out sample of statistically independent data.

#### 3.1 Fitting the Data

The notion of fitness 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 fitness for a set of machines is similar; a set of machine fits a

---

<sup>3</sup> See Engle-Warnick and Slonim (1999) for a formal presentation of the reduction process. See Hopcroft and Ullman (1979) and Harrison (1964) for finite automata theory. In addition to eliminating behaviorally equivalent machines, we examine each machine in our candidate set to determine if there are any behaviorally equivalent machines with fewer states. If there are, we replace the existing machine with the one with the fewest states. And if there are multiple machines with minimal states, we arbitrarily choose one.

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 fitness (we loosely refer to the 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 fitness 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 fitness  $F(M_i)$  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 time and players, we examine the fitness 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 fitness  $F(T_n)$  of set  $T_n$  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 fitness over every possible subset  $T_n$  of  $n$  machines in  $M$ . This is equivalent to minimizing the sum of squared errors:

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



Note that  $F(B_{n+1}) \geq F(B_n)$ , since  $B_{n+1}$  could always contain all the machines in  $B_n$ .<sup>4</sup> Thus the fitness 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 fitness 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 fitness 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 50 observations would need to be fit to increase the best fitting set size by 1). 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.<sup>5</sup>

### 3.2 Refining the Model Selection: Out of Sample Fitness

As a second guard against over fitting, we infer the best fitting set in one sample (the *training sample*) and test its fitness on an independent sample (*the test sample*).<sup>6</sup> We use a holdout sample to address the concern that when subjects play many

---

<sup>4</sup> Note also that  $B_n$  is not necessarily a subset of  $B_{n+1}$ .

<sup>5</sup> 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.

<sup>6</sup> For applications in pattern recognition see Devroye, Györfi, and Lugosi (1998).

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 fitness 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 fitness 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 fitness  $U(T_n, M_i)$  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 fitness 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 fitness 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 fitness 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.

## 4 The Trust Game Experiment

The experimental design is integral to determining what strategies people play in repeated games. We specifically chose the trust game and its parameters to induce actions to vary across rounds and to depend on opponent's choices. 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.

We chose an extensive form game to explore the robustness of the inference method to different degrees of information. And we chose a game that may induce behavior (e.g., trust and reciprocity) that strategy inference can offer novel insights into.

Figure 1 shows the extensive form game.<sup>7</sup> 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.

#### 4.1 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.<sup>8</sup> Each session began with 20 indefinite supergames and finished with 20 finite supergames.

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 14 rounds, and the shortest was 1 round. We ran two sessions with each sequence of supergame lengths and used one for the training sample and the other for the

---

<sup>7</sup> This game is theoretically analyzed by Kreps (1990). Berg, Dickhaut, and McCabe (1994) examine behavior when a similar game is played without repetition. See also Clark and Sefton (2001) and Glaeser, Laibson, Scheinkman, and Soutter (2000) for trust experiments.

<sup>8</sup> Instructions are available from the authors. Selten and Stoeker (1986) examine supergame play in finitely repeated PD games. Van Huyck, Battalio, and Walters (1997) study indefinitely repeated games using a similar stage game, with a continuum of actions.

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. Different expected supergame lengths, across the finite or indefinite condition might cause variation in behavior between the treatments; we eliminated this source of noise by choosing the same number of expected rounds across treatments.

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.

Four sessions were run with 14 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.<sup>9</sup> The experiments were run at the University of Pittsburgh. Subjects were paid a \$5.00 participation fee plus their earnings from

---

<sup>9</sup> Due to a computer failure we lost 7 indefinite observations (one round) and 98 finite observations in one session (14 rounds). 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.

four supergames that were randomly selected at the end of the session.

## 4.2 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 equilibrium behavior for both players in every round of the finite game. For the indefinite game there are many equilibria.<sup>10</sup> 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 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 in the experiment. If the discount factor (i.e., the continuation probability) is less than 0.75, then 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.

## 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 fitness 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 examines strategies inferred

---

<sup>10</sup> This is a result from the Folk Theorem of Repeated Games, see Fudenberg and Maskin (1986).

over time.

## 5.1 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.<sup>11</sup> 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  $M^B$ , 12 and 26, 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.

Our remaining criteria are based on theory, past experimental evidence and conjectures derived from protocols. Theory 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

---

<sup>11</sup> 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.

unravels from cooperative behavior towards the non-cooperative equilibrium.<sup>12</sup> To examine this unraveling 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.<sup>13</sup>

Protocols collected in pilot sessions of the indefinite game 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 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 unraveling 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  $+$ .

## 5.2 Fitness of the Candidate Sets

The top of each subsection of Table 1 shows the fitness 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 fitness 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 1 provides motivation for examining the  $2+$  candidate set. Note that while

---

<sup>12</sup> 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  $S$  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.

<sup>13</sup> For these strategies we assume during the cooperative faze that if an opponent plays his non-cooperative action then the player responds by playing his non-cooperative action thereafter.

the number of strategies in  $M^A$  increases by a factor of 10 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 30 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 fitness enough to risk over fitting the data with the 10 and 30 fold increases in candidate strategies.

### 5.3 Number of Strategies in the Data

The bottom of each subsection of Table 1 shows the fitness 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 fitness,  $F(B_n)$ ; holding  $n$  constant, the difference in fitness 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 1 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.<sup>14</sup> Inferring more Player B than A strategies may be because Player B has more behavioral options; e.g., Player B can exploit Player A (by playing  $K$ ), but Player A cannot exploit Player B. In general, Player B has more options since he always observes Player A actions whereas Player A does not always observe Player B actions.

---

<sup>14</sup> Computer simulations replicating game conditions indicate that  $c = 5\%$  is conservative in the sense that over fitting the data did not occur (see footnote 5 above). The simulations also validated the ability to recover the known data generating machines in the presence of noise.



The results in Table 1 are based on the training sample. 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 fitness 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.

#### 5.4 Specific Strategies in the Data

Figure 6 shows the inferred strategies and Table 2 shows their unique and total fitness in the training and test samples. Henceforth, we refer to each machine’s fitness (i.e., the number of machines that it fits, recall Definition 1) as its total fitness in contrast to its unique fitness (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 fitness.

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 two unconditional cooperation 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 although we did not *a priori* specify the model to look for possible individual or learning effects.

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 cooperation strategy ( $M_{a1}$ ) is a best response to Player B’s unconditional cooperation 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 ( $M_{a14}$ ).

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.<sup>15</sup> Thus,  $M_{a15}$  reflects a response to Player B behavior that was not inferred. And although Player B’s unconditional cooperation 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)). Thus, every inferred finite strategy is understandable; four are best responses to inferred opponent strategies, one is a best response to an action we did not infer and one reflects a utility function with preferences for fairness.

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.

---

<sup>15</sup> All best fitting strategy sets for Player B with more than 3 strategies include the 5th round defection strategy  $M_{b29}$ .

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). Finally, two of the inferred strategies in the indefinite game form a cooperative equilibrium pair.

## 5.5 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 10 supergames. Figure's 7 and 8 show the inferred finite and indefinite game strategies, respectively, for the first and last 10 supergames and Table 3 shows the fitness of the inferred strategies.

In the finite game, the same strategy sets that are inferred for all 20 supergames for Player A and Player B are inferred for the first 10 supergames (see Figure 7a for Player A and Figure 7b for Player B). However, for Player A for the last 10 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 10 supergames are replaced with two new strategies for the last 10 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 replacing the three and four round conditional counters ( $M_{b27}$  and  $M_{b28}$ ) with a

strategy ( $M_{b7}$ ) that plays  $R$  the first time Player A plays  $S$ , no matter which round Player A plays  $S$  the first time, and then plays  $K$  any time Player A plays  $S$  thereafter.

In the indefinite games (see Figure 8a for Player A and Figure 8b for Player B) we continue to only infer the Grim Trigger strategy for Player A ( $M_{a3}$ ). Over time Player A increasingly makes choices consistent with this strategy; its fitness rises from an average of 65% in the first 10 to 74% in the last 10 supergames. Likewise, Player B increasingly makes choices consistent with machine  $M_{b6}$ , fitting 44% in the first 10 and 58% in the last 10 supergames. Thus, the inferred equilibria pair ( $M_{a3}$ ,  $M_{b6}$ ) over all 20 supergames is inferred in both the first and last 10 supergames and 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 10 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 20 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  $K$ .<sup>16</sup>

In sum, behavior in the finite and indefinite games evolves in a logical manner. In the indefinite game, both players increasingly play strategies consistent with a

---

<sup>16</sup> We may infer different behavior for Player B after he plays  $K$  across the first and last ten supergames as well as across all 20 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. The lack of observations to distinguish between strategies because certain paths of play rarely occur is addressed in Engle-Warnick and Slonim (1999). To address this concern, they proceed as follows. First, they find, for each inferred machine, sets of machines that contribute to the best fitting set's fitness as well, or almost as well, as the inferred machine. Second, they find the common component behaviors of the machines within these sets, and conclude that the common components are the important behaviors to fit the data. This procedure guards against attributing too much meaning to a strategy if other strategies could have fit the data almost as well. For example, there are many similar fitting machines to the inferred Player B machines that play  $K$  in round 1, and the only common component of these machines is in fact to play Keep in round 1.

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 20 supergames may be partially explained by players adapting different strategies over time.

## 6 Conclusion

This paper combines a strategy model, an inference procedure and experimental observations to examine unobserved repeated game strategies from observed actions. Inference on specific repeated-game strategies is vital to bridge the gap between theory and behavior because it allows researchers to form and examine hypotheses regarding strategies based on empirical observation.

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 the trigger strategy is harsh enough that it may be included in the construction of repeated-game equilibria. Further, in the finite and indefinite games the inferred Player A and B strategies are best responses to the inferred strategies of their opponent and the inferred strategies evolve in a logical (best response) manner.

For both types of repeated games and both types of players we use the same inference procedure and candidate strategy sets. To the extent that the same procedure produces distinct but intuitive and theoretically justified results across the two conditions, the inference method is robust.

We simplified the environment in several ways to overcome the difficulty that strategies are unobserved. For example, we limited the stage-game strategy space to two actions for each player and we examined relatively short supergame lengths to limit noise in the data-generating process. To distinguish between different strategies, we collected over 500 supergame decisions in order to observe many different realizations of actions. To safeguard against over fitting, we introduced a cost func-

tion that penalized inferred strategy sets proportionally to the number of strategies in the set, and checked out of sample predicted fitness of inferred strategies. The out of sample test also guards against spuriously fitting data due to dependencies of many observations per subject and possible path dependencies of play. To make the results easily interpretable we examined deterministic strategies and used finite automata to model the decision rules players may use.

Once strategies are identified, they can be used to address many important questions. For example, what are the intermediate and long-term convergence properties of games in which players use the inferred strategies? Theory and computational models of learning can address these questions (i.e., the method can be used in conjunction with any model requiring strategy inference (e.g., Camerer and Ho (1999); Roth and Erev (1995)). The method can also be used to study other repeated games (e.g., public goods and social dilemma problems).

This paper's goal is to take a step toward inferring repeated-game strategies, a core part of game theory, in natural environments. We believe that a simple strategy model like the one used in this paper can be extended to the field since there are many instances where limiting the action space to two or three actions may be reasonable. Examples include union decisions regarding whether to go on strike and Federal Reserve Board decisions regarding whether to increase, decrease or maintain interest rates. Experiments enable researchers to examine many instances of play of repeated games; no such luxury exists in the field. As such, extensions to field work may include relaxing the assumption of deterministic machines. The goal is to advance game theory as an empirical science.

## References

- ABREU, D., AND A. RUBINSTEIN (1988): “The Structure of Nash Equilibrium in Repeated Games With Finite Automata,” *Econometrica*, 56, 1259–1282.
- AXELROD, R. (1984): *The Evolution of Cooperation*. Harper Collins, New York.
- BERG, J., J. DICKHAUT, AND K. MCCABE (1994): “Trust, Reciprocity, and Social History,” *Games and Economic Behavior*, 10, 278–305.
- BINMORE, K., AND L. SAMUELSON (1992): “Evolutionary Stability in Repeated Games Played by Finite Automata,” *Journal of Economic Theory*, 57, 278–305.
- BOLTON, G., AND A. OCKENFELS (2000): “ERC: A Theory of Equity, Reciprocity and Competition,” *The American Economic Review*, 90, 166–193.
- BRANDTS, J., AND G. CHARNES (2000): “Hot vs. Cold: Sequential Responses and Preference Stability in Experimental Games,” *Experimental Economics*, 2, 227–238.
- CAMERER, C., AND T. HO (1999): “Experience-weighted Attraction (EWA) Learning in Normal-form Games,” *Econometrica*, 67, 827–874.
- CLARK, K., AND M. SEFTON (2001): “The Sequential Prisoner’s Dilemma: Evidence on Reciprocation,” *The Economic Journal*, 111, 51–68.
- COSTA-GOMEZ, M., V. CRAWFORD, AND B. BROSETA (2001): “Cognition and Behavior in Normal-Form Games: An Experimental Study,” *Econometrica*, Forthcoming.
- DEVROYE, L., L. GYORFI, AND G. LUGOSI (1998): *A Probabilistic Theory of Pattern Recognition*. Springer-Verlag, New York.
- EL-GAMAL, M., AND D. GREETHER (1995): “Are People Bayesian? Uncovering Behavioral Strategies,” *Journal of the American Statistical Association*, 90, 1137–1145.
- ENGLE-WARNICK, J. (2001): “Inferring Strategies from Experimental Data: A Nonparametric, Tree Classification Approach,” Unpublished manuscript.
- ENGLE-WARNICK, J., AND R. SLONIM (1999): “Inferring Repeated Game Strategies From Actions: Evidence From a Trust Game Experiment,” Unpublished manuscript.
- FEHR, E., AND K. SCHMIDT (1999): “A Theory of Fairness, Competition, and Cooperation,” *The Quarterly Journal of Economics*, 114, 817–868.
- FUDENBERG, D., AND E. MASKIN (1986): “The Folk Theorem of Repeated Games With Discounting or With Incomplete Information,” *Econometrica*, 54, 533–554.
- GLAESER, E., D. LAIBSON, J. SCHEINKMAN, AND C. SOUTTER (2000): “Measuring Trust,” *The Quarterly Journal of Economics*, 125, 811–846.

- HARRISON, M. (1964): *Introduction to Switching and Finite Automata*. Addison-Wesley, Reading, MA.
- HOPCROFT, J., AND J. ULLMAN (1979): *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, Reading, MA.
- KALAI, E., AND W. STANFORD (1988): "Finite Rationality and Interpersonal Complexity in Repeated Games," *Econometrica*, 565, 397–410.
- KREPS, D. (1990): "Corporate Culture and Economic Theory," in *Perspectives on Positive Political Economy*, ed. by J. Alt, and K. Shepsle, pp. 90–143. Cambridge University Press, Cambridge.
- MILLER, J. (1996): "The Coevolution of Automata in the Repeated Prisoner's Dilemma," *Journal of Economic Behavior and Organization*, 29, 87–112.
- MOORE, E. (1956): "Gedanken Experiments on Sequential Machines," in *Automata Studies*. Princeton University Press, Princeton.
- OSBORNE, M., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. MIT Press, Cambridge, MA.
- OSTROM, E. (1998): "A Behavioral Approach to the Rational Choice Theory of Collective Action: Presidential Address, American Political Science Association," *American Political Science Review*, 92, 1–22.
- ROTH, A., AND I. EREV (1995): "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior, Special Issue, Nobel Symposium*, 8, 164–212.
- SELTEN, R., AND M. MITZKEWITZ (1997): "Duopoly Strategies Programmed by Experienced Traders," *Econometrica*, 65, 517–555.
- SELTEN, R., AND R. STOEKER (1986): "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames," *Journal of Economic Behavior and Organization*, 65, 47–70.
- SLONIM, R. (1994): "Learning and Bounded Rationality in a Search-for-the-best-alternative Experiment," *Journal of Economic Behavior and Organization*, 25, 141–165.
- (1999): "Learning Rules of Thumb or Learning More Rational Rules," *Journal of Economic Behavior and Organization*, 38, 217–236.
- STAHL, D. (2000): "Rule Learning in Symmetric Normal Form Games: Theory and Evidence," *Games and Economic Behavior*, 32, 105–138.
- VAN HUYCK, J., R. BATTALIO, AND M. WALTERS (1997): "Is Reputation a Substitute for Commitment in the Peasant-dictator Game?," Unpublished manuscript.



**Table 1: Fitness of All and Best Fitting Machines in the Training Sample**

| FINITE GAMES   |           |     |     |       |     |     |  |           |      |     |       |     |     |
|--|-----------|-----|-----|-------|-----|-----|--|-----------|------|-----|-------|-----|-----|
| A Players  |           |     |     |       |     |     | B Players                                  |           |      |     |       |     |     |
| No of Observations O: 280  |           |     |     |       |     |     | *No of Observations O: 275                 |           |      |     |       |     |     |
| No. of Machines  |           |     |     |       |     |     | No. of Machines                            |           |      |     |       |     |     |
| Fitness  |           |     |     |       |     |     | Fitness                                    |           |      |     |       |     |     |
| % of O   |           |     |     |       |     |     | % of O                                     |           |      |     |       |     |     |
| All  | 1 state   | 2   | 70  | 25.0% |     |     | All  | 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 | 75.6% |     |     |
|  | 1 state + | 10  | 255 | 91.1% |     |     |  | 1 state + | 10   | 265 | 96.4% |     |     |
|  | 2 state + | 18  | 256 | 91.4% |     |     |  | 2 state + | 32   | 268 | 97.5% |     |     |
|  | 3 state + | 180 | 278 | 99.3% |     |     |  | 3 state + | 1058 | 274 | 99.6% |     |     |
| No. in: Fitness of Best Fitting Sets from:   |           |     |     |       |     |     | No. in: Fitness of Best Fitting Sets from: |           |      |     |       |     |     |
| B <sub>n</sub>   | 1         | 2   | 3   | 1+    | 2+  | 3+  | B <sub>n</sub>                             | 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 |
| *Note - Five observations are missing since the A Player did not send for five entire supergames |           |     |     |       |     |     |  |           |      |     |       |     |     |
| INDEFINITE GAMES   |           |     |     |       |     |     |  |           |      |     |       |     |     |
| A Players  |           |     |     |       |     |     | B Players                                  |           |      |     |       |     |     |
| No of Observations O: 280  |           |     |     |       |     |     | No of Observations O: 280                  |           |      |     |       |     |     |
| No. of Machines  |           |     |     |       |     |     | No. of Machines                            |           |      |     |       |     |     |
| Fitness  |           |     |     |       |     |     | Fitness                                    |           |      |     |       |     |     |
| % of O   |           |     |     |       |     |     | % of O                                     |           |      |     |       |     |     |
| All  | 1 state   | 2   | 170 | 60.7% |     |     | All  | 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% |     |     |
| No. in: Fitness of Best Fitting Sets from:   |           |     |     |       |     |     | No. in: Fitness of Best Fitting Sets from: |           |      |     |       |     |     |
| B <sub>n</sub>   | 1         | 2   | 3   | 1+    | 2+  | 3+  | B <sub>n</sub>                             | 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 |

**Table 2: Inferred Strategies and Their Fitness**

|                 | <u>Machine</u> | <u>Training Sample</u> |                      | <u>Test Sample</u>    |                      |
|-----------------|----------------|------------------------|----------------------|-----------------------|----------------------|
|                 |                | <u>Unique Fitness</u>  | <u>Total Fitness</u> | <u>Unique Fitness</u> | <u>Total Fitness</u> |
| Finite Player A | $M_{a15}$      | 20%                    | 40%                  | 15%                   | 42%                  |
|                 | $M_{a14}$      | 17%                    | 38%                  | 8%                    | 35%                  |
|                 | $M_{a1}$       | 23%                    | 23%                  | 9%                    | 9%                   |
| Finite Player B | $M_{b1}$       | 20%                    | 44%                  | 12%                   | 29%                  |
|                 | $M_{b28}$      | 15%                    | 41%                  | 13%                   | 43%                  |
|                 | $M_{b27}$      | 10%                    | 19%                  | 17%                   | 38%                  |
| Indef. Player A | $M_{a3}$       | 72%                    | 72%                  | 67%                   | 67%                  |
| Indef. Player B | $M_{b6}$       | 43%                    | 57%                  | 32%                   | 46%                  |
|                 | $M_{b11}$      | 7%                     | 22%                  | 5%                    | 21%                  |
|                 | $M_{b25}$      | 11%                    | 11%                  | 15%                   | 17%                  |

1. Based on  $c = 5\%$  threshold

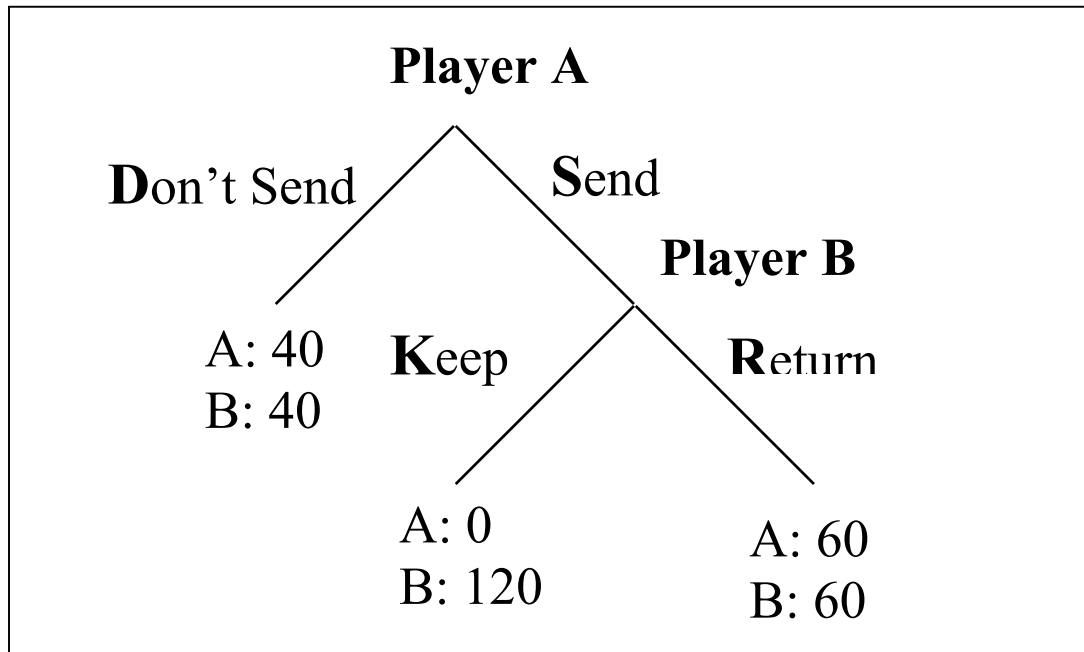
**Table 3: Inferred Strategies and Their Fitness Over Time**

|                 | FIRST TEN SUPERGAMES |                        |                      |                       |                      | LAST TEN SUPERGAMES |                        |                      |                       |                      |
|-----------------|----------------------|------------------------|----------------------|-----------------------|----------------------|---------------------|------------------------|----------------------|-----------------------|----------------------|
|                 | <u>Machine</u>       | <u>Training Sample</u> |                      | <u>Test Sample</u>    |                      | <u>Machine</u>      | <u>Training Sample</u> |                      | <u>Test Sample</u>    |                      |
|                 |                      | <u>Unique Fitness</u>  | <u>Total Fitness</u> | <u>Unique Fitness</u> | <u>Total Fitness</u> |                     | <u>Unique Fitness</u>  | <u>Total Fitness</u> | <u>Unique Fitness</u> | <u>Total Fitness</u> |
| Finite Player A | M <sub>a15</sub>     | 28%                    | 42%                  | 21%                   | 47%                  | M <sub>a15</sub>    | 46%                    | 46%                  | 34%                   | 34%                  |
|                 | M <sub>a14</sub>     | 15%                    | 29%                  | 10%                   | 36%                  | M <sub>a5</sub>     | 27%                    | 27%                  | 6%                    | 6%                   |
|                 | M <sub>a1</sub>      | 26%                    | 26%                  | 13%                   | 13%                  |                     |                        |                      |                       |                      |
| Finite Player B | M <sub>b1</sub>      | 24%                    | 43%                  | 15%                   | 31%                  | M <sub>b30</sub>    | 41%                    | 46%                  | 25%                   | 34%                  |
|                 | M <sub>b28</sub>     | 20%                    | 42%                  | 18%                   | 44%                  | M <sub>b7</sub>     | 17%                    | 21%                  | 7%                    | 15%                  |
|                 | M <sub>b27</sub>     | 11%                    | 17%                  | 15%                   | 30%                  |                     |                        |                      |                       |                      |
| Indef. Player A | M <sub>a3</sub>      | 67%                    | 67%                  | 62%                   | 62%                  | M <sub>a3</sub>     | 77%                    | 77%                  | 71%                   | 71%                  |
| Indef. Player B | M <sub>b6</sub>      | 54%                    | 54%                  | 62%                   | 62%                  | M <sub>b6</sub>     | 28%                    | 61%                  | 17%                   | 55%                  |
|                 | M <sub>b16</sub>     | 14%                    | 14%                  | 20%                   | 20%                  | M <sub>b24</sub>    | 9%                     | 9%                   | 13%                   | 15%                  |
|                 |                      |                        |                      |                       |                      | M <sub>b28</sub>    | 9%                     | 42%                  | 11%                   | 51%                  |


1. Based on  $c = 5\%$  threshold

2. Arrows  $\rightarrow$  reflect similar machines that they have almost identical structure except for one or two transitions

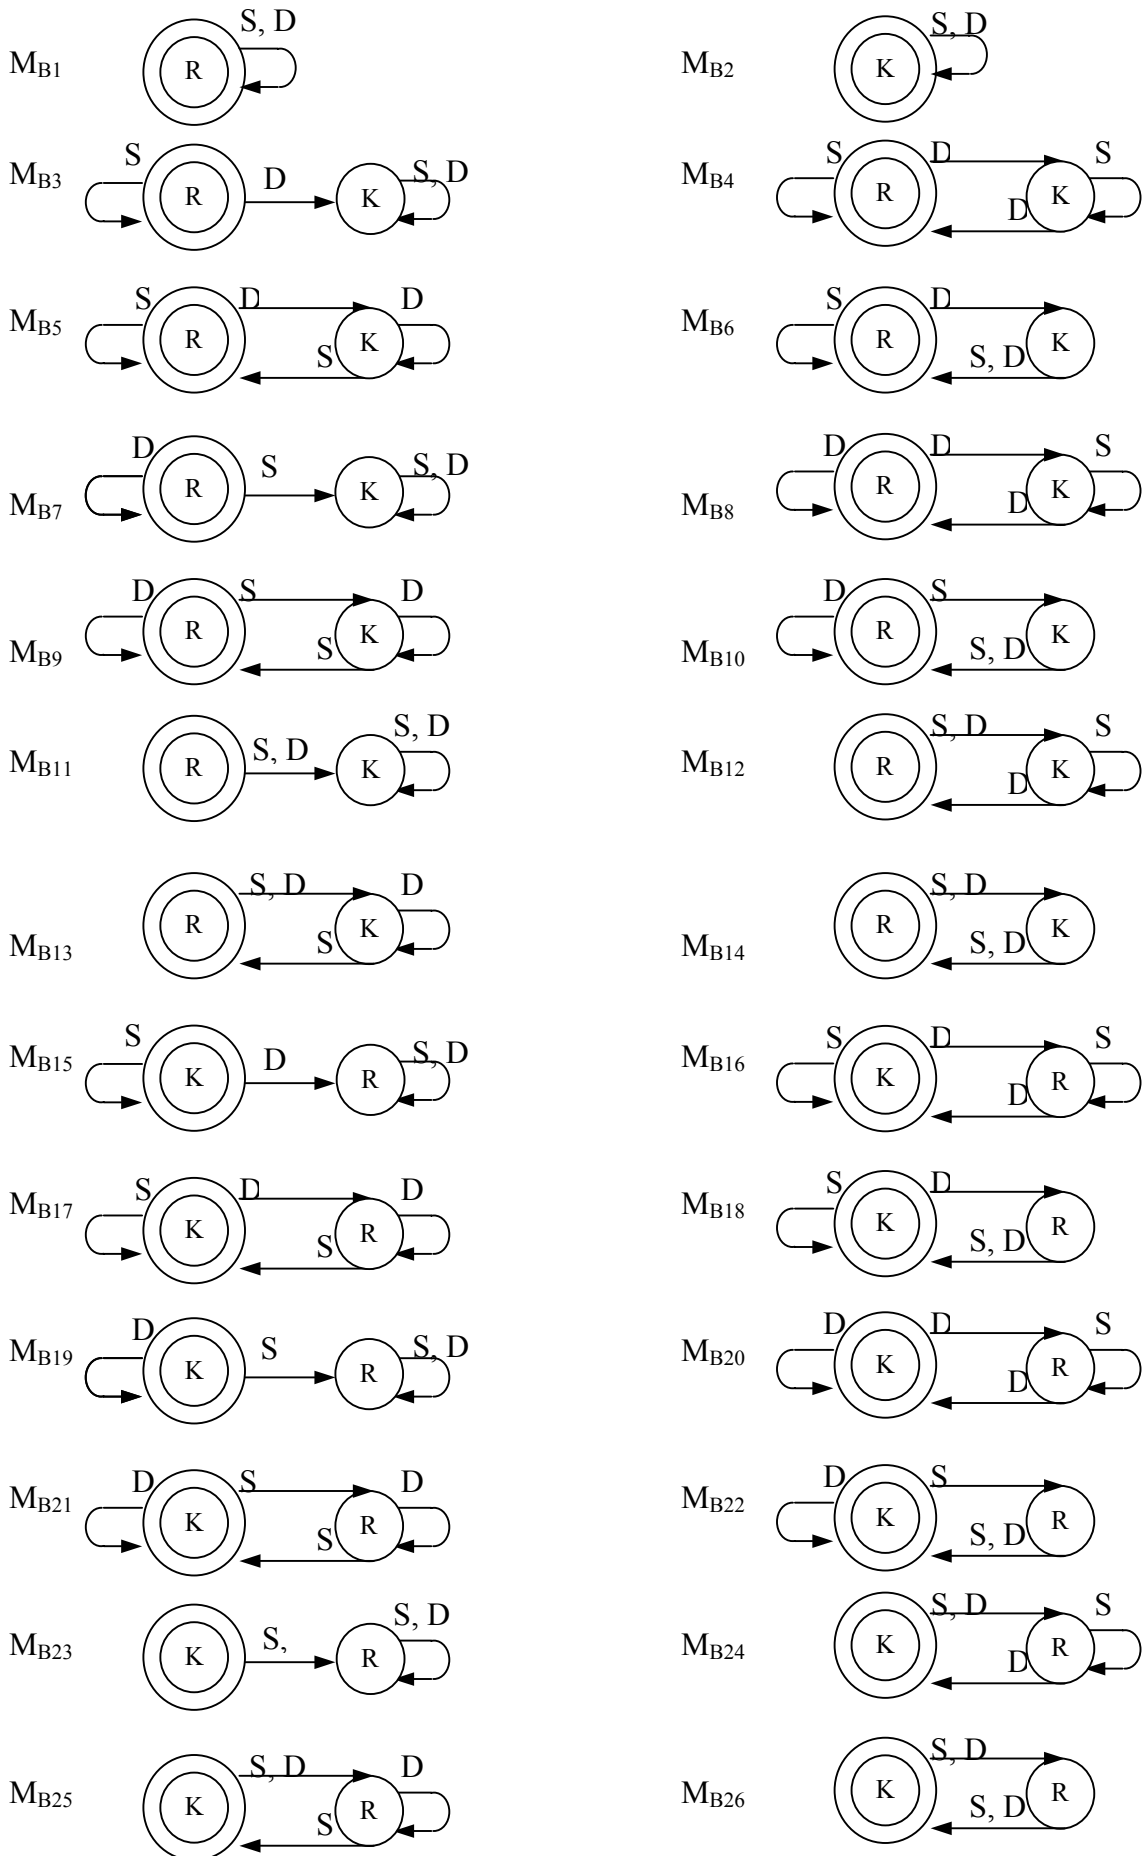
Figure 1: Trust Game Extensive Form




**Figure 2: Player A, Set  $s \leq 2$**




**Figure 3: Player B, Set  $s \leq 2$**



**Figure 4: Player A, Set “+”**




**End Game Effect Strategies**

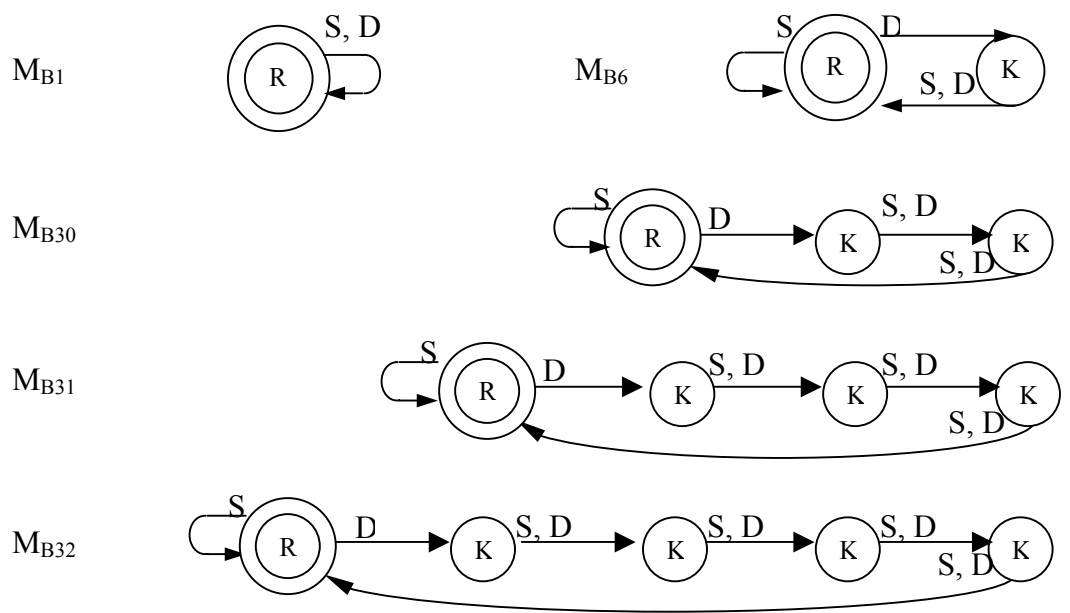


**Punishment Strategies**

**Figure 5: Player B, Set “+”**



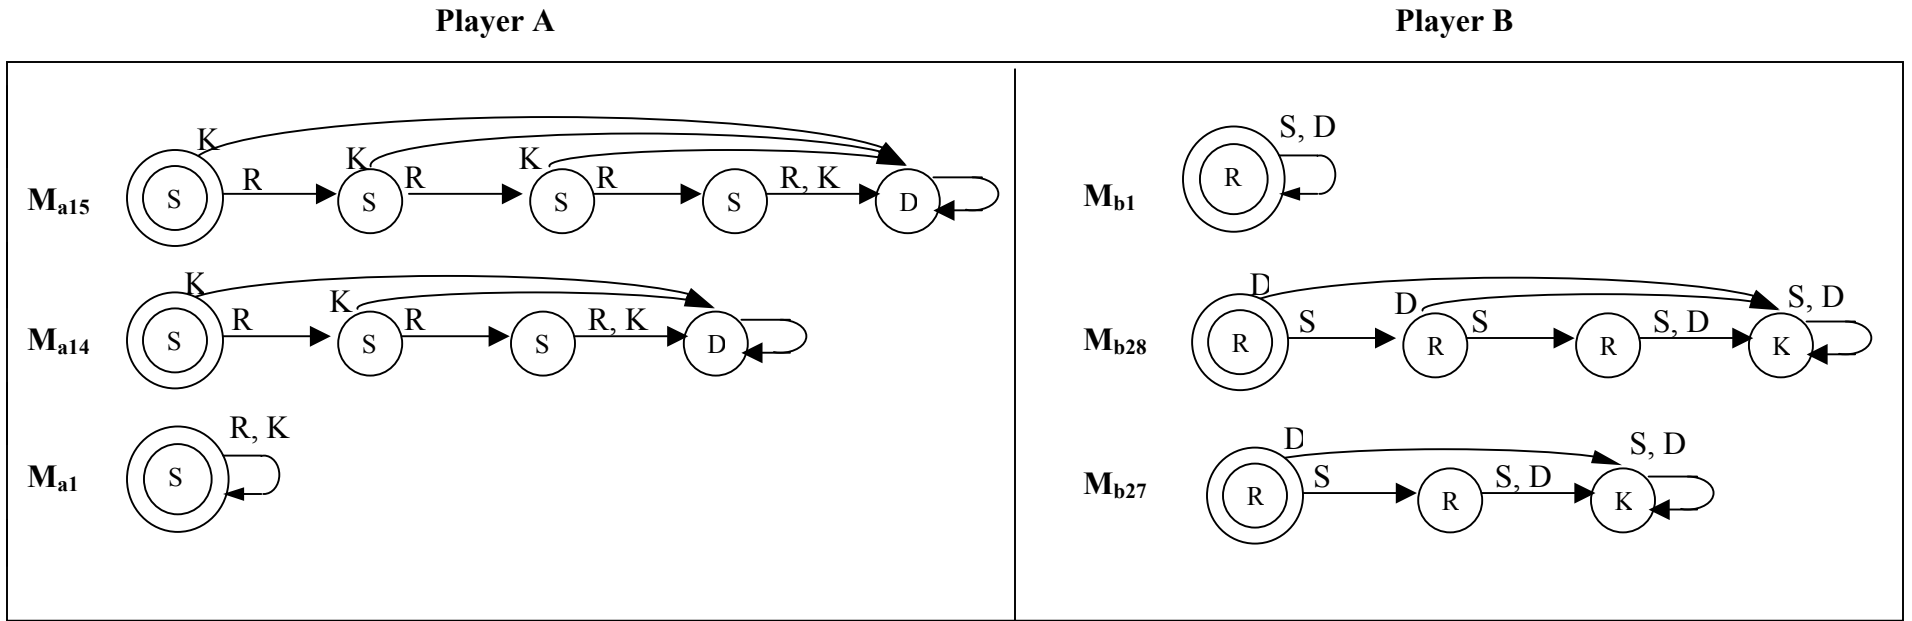
**End Game Effect Strategies**




**Punishment Strategies**



**Figure 6a: Inferred Strategies In All Finite Supergames**



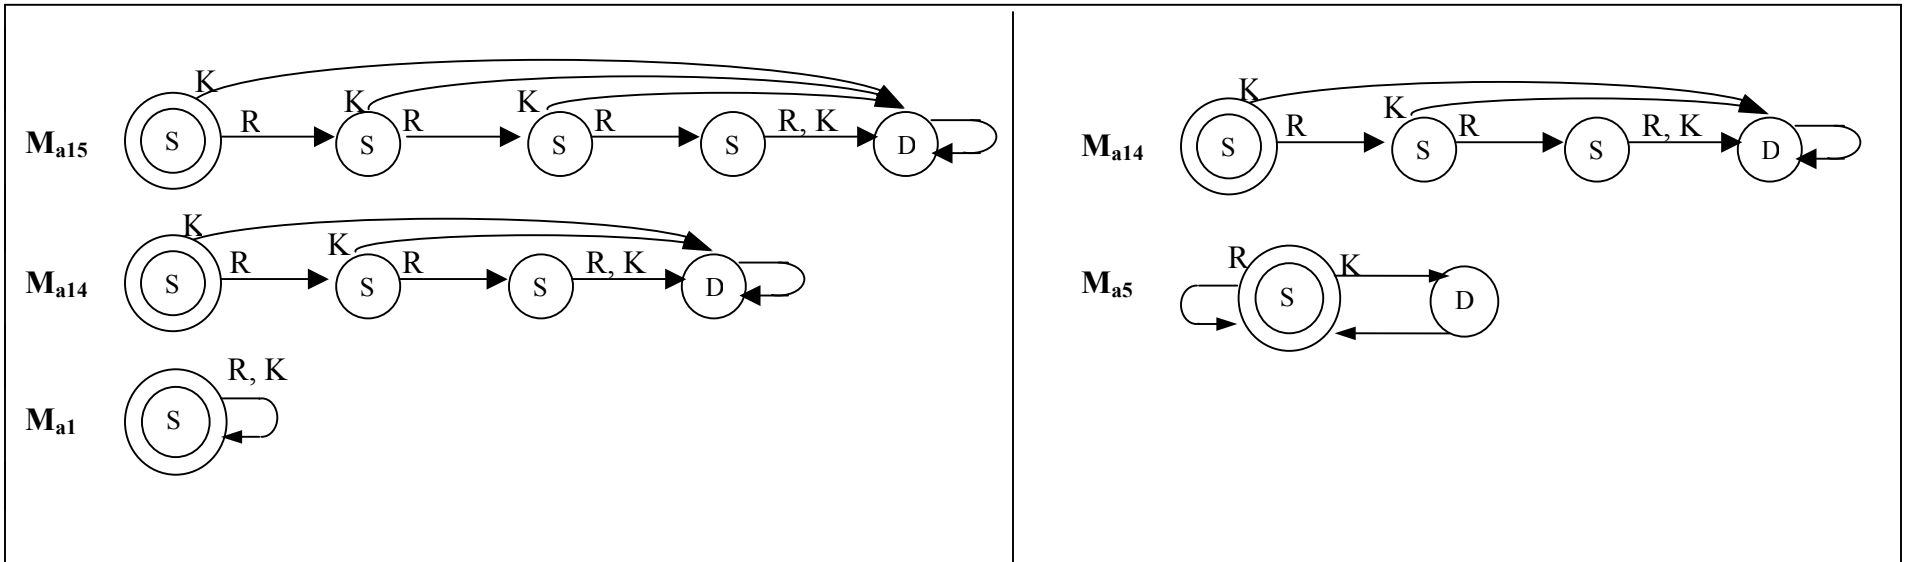
**Figure 6b: Inferred Strategies In All Indefinite Supergames**



**Figure 7a: Inferred Player A Strategies In Finite Supergames**

**Player A: First 10 Supergames**


**Player A: Last 10 Supergames**



**Figure 7b: Inferred Player B Strategies In Finite Supergames**

**Player B: First 10 Supergames**


**Player B: Last 10 Supergames**



**Figure 8a: Inferred Player A Strategies In Indefinite**

Player A: First 10 Supergames

Player A: Last 10 Supergames



**Figure 8b: Inferred Player B Strategies In Indefinite**

Player B: First 10 Supergames

Player B: Last 10 Supergames

