

Implementation by Iterative Dominance and Backward Induction: An Experimental Comparison¹

Elena Katok

*Smeal College of Business Administration, Penn State University,
University Park, Pennsylvania 16802*

ekatok@psu.edu

Martin Sefton

School of Economics, University of Nottingham, Nottingham NG7 2RD, United Kingdom

martin.sefton@nottingham.ac.uk

and

Abdullah Yavas

*Smeal College of Business Administration, Penn State University,
University Park, Pennsylvania 16802*

ayavas@psu.edu

Received January 18, 2001; final version received February 16, 2001;
published online December 4, 2001

We report experimental results on the relative performance of simultaneous and sequential versions of the Abreu–Matsushima mechanism. Under the simultaneous version, subjects typically use undominated strategies, but apply only a limited number of iterations of dominance. Consequently the unique strategy surviving iterative elimination of strictly dominated strategies is rarely observed. Under the sequential version, subjects also typically use undominated strategies, but apply only a limited number of steps of backward induction. Thus the backward induction outcome is also rarely observed. The sequential version results in fewer observed outcomes corresponding to the predicted outcome than the simultaneous version. *Journal of Economic Literature* Classification Number(s): C72, C92. © 2001 Elsevier Science (USA)

¹ We thank Geir Asheim, John Morgan, Motty Perry, Daniel Seidmann, and an associate editor for useful comments. Funding for the experiments from the Smeal College of Business Administration and the Uni-Marts Endowment at Penn State University is gratefully acknowledged.

1. INTRODUCTION

In this paper we consider two mechanisms that can, in principle, be applied to a wide range of games to implement desired outcomes, and compare their performances in a laboratory environment. The first mechanism, proposed by Abreu and Matsushima [1], hereafter AM, virtually implements almost any social choice function via iterative elimination of strictly dominated strategies. The second mechanism, proposed by Glazer and Perry [7], hereafter GP, is a modification of the AM mechanism that implements via backward induction.

An appealing feature of iterative strict dominance is that it can be based on simple rationality postulates. If players use undominated strategies and this is common knowledge, then the outcome of a game will be that selected by iterative strict dominance. However, some authors have argued that when many iterations of dominance are required, and each involves a small payoff difference, this solution concept is not compelling. Glazer and Rosenthal [8] argue that the AM mechanism implements outcomes in a manner that is susceptible to this criticism (but see also AM's reply, [2]).

In a laboratory experiment, Sefton and Yavas [15] investigated the AM mechanism and found that the game-theoretic prediction is indeed a poor predictor of subjects' behavior. While subjects overwhelmingly do play undominated strategies, at the same time they fail to drive the logic of iterative dominance to its conclusion. Rather, subjects can be loosely characterized as only carrying out a limited number of applications of iterated dominance (usually 3 or 4). In fact, limited application of iterative dominance is typical of experimental results from other games. For example Costa-Gomes *et al.* [6] find that compliance with equilibrium is high in normal-form games that are solvable by one or two rounds of iterated dominance, but much lower in games that require three rounds. In dominance-solvable "beauty contest game" experiments there is also evidence that subjects apply a limited number of rounds of dominance, although they tend to apply more rounds with experience: see Nagel [12] for a survey.

These results suggest that the AM mechanism may not be robust to one or more of a number of possible factors—lack of common knowledge, bounded rationality, or uncontrolled aspects of subjects preferences. Thus, it is natural to ask whether other mechanisms provide robust alternatives to this mechanism. Glazer and Rubinstein [9] point out that it is sometimes possible to construct an extensive game representation of a normal form game with the property that the backward induction outcome of the extensive game involves the same comparisons between strategies as is made in a particular order of elimination of strategies in the normal form

game. In such a case they argue that the calculation of the backward induction outcome may be simpler than the calculation of the outcome surviving iterative dominance, since the extensive form provides an explicit guide to the order in which strategies should be compared. The GP mechanism can be viewed as such an extensive game. This mechanism is a dynamic version of AM that implements the desired outcome via backward induction. GP argue that sequential mechanisms, relying on backward induction, seem to be more intuitive and simpler to understand than simultaneous counterparts.

However, as with iterative dominance, some authors have stated reservations about the backward induction solution concept. For example, Selten [16] argues that the backward induction outcome is an implausible prediction in a game where there are many subgames to consider, and Rosenthal [13] points out that the effect of decision costs or a small probability of a “mistake” may be compounded in games with many subgames.² There is also evidence of the limited predictive power of backward induction in experimental games. In experiments with centipede games subjects’ decisions do not lead to the backward induction outcome, although decisions move closer to the prediction as subjects repeat the game (McKelvey and Palfrey [11]). In alternating offer bargaining games subjects also systematically deviate from the backward induction outcome (see Roth [14] for a survey), although for some parameter values outcomes closer to the backward induction can be obtained by giving subjects experience in the relevant subgames (Harrison and McCabe [10]). These experiments suggest that the vulnerability of iterative dominance to lack of common knowledge, bounded rationality, or uncontrolled preferences is shared, to some extent, by backward induction.

Which solution concept is more robust is difficult to determine using existing experimental evidence, because the games in which each has been studied differ considerably. In the present paper we compare the predictive performance of the AM and GP mechanisms using the game discussed by Glazer and Rosenthal [8].

The sessions involving the AM mechanism are based on Experiment I of Sefton and Yavas [15], but involve slightly different parameters and procedures. The results reported here are broadly consistent with those from the earlier study: subjects overwhelmingly play undominated strategies, but do not drive the logic of iterative dominance to its conclusion; as a result the predicted outcome is rarely observed. Sessions with the GP mechanism are designed to be comparable with these AM sessions. Again, the vast

²In addition, the dynamic structure introduces difficult issues involving what rational players should do when they find themselves off the backward induction path. For different approaches to these issues see Aumann [3] and Ben-Porath [4].

majority of subjects play undominated strategies, but this time the unravelling process that leads to the backward induction outcome fails; as a result the backward induction outcome is rarely observed. Indeed, we observe relatively fewer outcomes corresponding to the predicted outcome under the GP mechanism than under the AM mechanism.

The remainder of the paper is organized as follows. In the next section we describe the games studied, and in Section 3 we describe our experimental design and procedures. In Section 4 we present the results and Section 5 concludes.

2. THE MECHANISM GAME

The games investigated here are based upon the coordination game of Fig. 1. (This game was used by Glazer and Rosenthal [8] as an example to illustrate some of their concerns with the way that the AM mechanism works.) The game is a symmetric two-player game with two pure-strategy equilibria, (α, α) and (β, β) .

We compare the performances of the AM and GP mechanisms in implementing the outcome (α, α) . Both mechanisms modify the game in two ways. First, players play the game in *pieces*. Instead of simultaneously making a single choice, each player chooses a sequence of T choices, where each choice is either α or β . These sequences are then matched choice by choice, and for each pair of choices payoffs are those given in Fig. 1 divided by T . Second, a player who includes a choice of β in his or her sequence may have to pay a small *fine*, F .

2.1. The Abreu–Matsushima Mechanism

In the AM mechanism the sequences are chosen simultaneously. A fine is then imposed on the player whose earliest choice of β in her sequence occurs before that of her opponent (they both incur the fine if their earliest choices of β occur in the same stage).

| | | |
|----------|----------|----------|
| | α | β |
| α | 100, 100 | 0, 0 |
| β | 0, 0 | 200, 200 |

FIG. 1. Payoff matrix for coordination game.

Consider the following example with $T = 10$ pieces and a fine of $F = 50$. Players one and two submit the following sequences:

| stage | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|------------|----------|---------|---------|---------|---------|---------|---------|---------|---------|---------|
| Player one | β | β | β | β | β | β | β | β | β | β |
| Player two | α | β |

For the first stage each player receives a payoff of zero, and for the other nine stages each player receives a payoff of 20, so the payoffs from pieced play are 180 to each player. Since player one's earliest choice of β occurs at an earlier stage (stage 1) than player two's (stage 2), player one incurs a fine of 50. Thus player one receives a payoff of 130, and player two receives a payoff of 180.

Formally, this is a symmetric game with two players, indexed by i , $i = 1, 2$. Each player's set of actions consists of T -tuples $x_i = (x_{i1}, \dots, x_{iT})$, where $x_{ik} \in \{0, 1\}$ (i.e., we denote α by 0 and β by 1). Player i 's payoff is given by

$$\pi_i = \sum_{k=1}^T \left(x_{ik} x_{jk} \frac{200}{T} + (1 - x_{ik})(1 - x_{jk}) \frac{100}{T} \right) - \sum_{k=1}^T (x_{ik} A_{ik} F),$$

where

$$A_{i1} = 1$$

$$A_{ik} = A_{i, k-1} (1 - x_{i, k-1}) (1 - x_{j, k-1}), \quad k = 2, \dots, T.$$

The first summation in the payoff function is the payoff from pieced play, while the second incorporates the fining part of the mechanism. The fining part of the mechanism depends on a recursive relation for A_{ik} . When $A_{ik} = 1$ the fining term is "switched on" and i incurs a fine by choosing β in the k th stage, and when $A_{ik} = 0$ the fining term is "switched off." The recursion works as follows: in the first stage the fining term is switched on, so either player incurs a fine if the first choice in their sequence is β . The fining term then remains switched on as long as both players choose α . As soon as a player chooses β , the fine is incurred by that player (possibly both) and the fining term is switched off.

When $F > 200/T$, the following results are easily established.

(1) Let x_{i2}, \dots, x_{iT} be arbitrarily given. The strategy $(0, x_{i2}, \dots, x_{iT})$ dominates $(1, x_{i2}, \dots, x_{iT})$.³

³ Throughout we focus on pure strategies. All dominance relations refer to strict dominance.

(2) Any strategy with $x_{i1} = 0$ is undominated.

(3) The unique strategy profile surviving iterative elimination of dominated strategies has $x_{ik} = 0$ for $i = 1, 2, k = 1, 2, \dots, T$.

To see (1), first suppose that player j 's strategy has $x_{j1} = 1$. Then player i 's payoff from the strategy beginning with $x_{i1} = 1$ is $P + 200/T - F$, where P represents her payoff from the last $T - 1$ pieces. Player i 's payoff from the same strategy but replacing $x_{i1} = 0$ is P . Thus, if $F > 200/T$ player i can avoid the fine by delaying her first choice of β , and this outweighs her first piece payoff loss of $200/T$. Alternatively, suppose player j 's strategy has $x_{j1} = 0$. Then the strategy beginning with $x_{i1} = 1$ gives player i a payoff of $P - F$, while the strategy beginning with $x_{i1} = 0$ yields at least $P + 100/T - F$. Thus, even if delaying the first choice of β does not lead to the avoidance of the fine, the increase in first piece payoff makes the choice $x_{i1} = 0$ worthwhile.

In fact, suppose j 's strategy is $x_j = (1, x_{j2}, \dots, x_{jT})$ where x_{j2}, \dots, x_{jT} are arbitrarily given. Then i 's strict best response is $x_i = (0, x_{j2}, \dots, x_{jT})$. Since a strategy cannot be dominated if it is a best response to some strategy, result (2) follows. Note that this means even a "non-monotonic" sequence such as $(0, 1, 0, \dots, 0)$ is undominated.⁴

Assuming common knowledge of rationality, property (1) implies that player i knows her opponent's sequence will begin with $x_{j1} = 0$. Knowing this, the same calculations as before imply that sequences with $x_{i2} = 1$ are strictly dominated by the same sequence but replacing $x_{i2} = 1$ with $x_{i2} = 0$. Continuing this process of iterative elimination of strictly dominated strategies leads to result (3).

2.2. The Glazer–Perry Mechanism

In the GP mechanism the players complete their choices sequentially: player 1 makes her first choice, this choice is then observed by player 2, player 2 then makes her first choice, and so on. Finally, having observed the T choices of her opponent, player 2 chooses the T th element of her own sequence (and thus the game is one of perfect information). The player who last chooses β incurs the fine (and since the players move sequentially only one player can incur the fine).

Consider the following example, again with $T = 10$ and $F = 50$:

| stage | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|------------|---------|---------|---------|---------|---------|---------|---------|---------|---------|----------|
| Player one | β |
| Player two | β | α |

⁴ Formally, monotonic sequences have the property that $k > m \Rightarrow x_{ik} \geq x_{im}$.

As in the earlier example, the payoffs from pieced play are 180 to each player. Here, player 1 made the last choice of β (in stage 10), and so incurs a fine of 50. Thus player 1 receives a payoff of 130, and player 2 receives a payoff of 180.

Thus, for the GP mechanism the recursive relation for A_{ik} in player i 's payoff function is replaced by

$$A_{2T} = 1$$

$$A_{1k} = A_{2k}(1 - x_{2k}), \quad k = 1, \dots, T$$

$$A_{2k} = A_{1,k+1}(1 - x_{1,k+1}), \quad k = 1, \dots, T-1.$$

Note that the GP mechanism introduces an asymmetry in player roles. For player 2, in the last stage the fining term is switched on so she incurs a fine if her last element is β . The fining term remains switched on as long as both players choose α as we work backwards through the sequences. As soon as a player chooses β (again, working backwards), the fine is incurred by that player and the fining term is switched off.

When $F > 200/T$, the following results are easily established:

(4) Let x_1 and $x_{21}, \dots, x_{2,T-1}$ be arbitrarily given. Then $x_{2T} = 0$ maximizes player 2's payoff.

(5) The unique backward induction outcome has $x_{ik} = 0$ for $i = 1, 2$, $k = 1, 2, \dots, T$.

To see (4), first suppose that player 1's last choice is β , i.e., $x_{1T} = 1$. Then player 2's payoff from choosing $x_{2T} = 1$ is $P + 200/T - F$, where P is the payoff from the first $T-1$ pieces, while her payoff from choosing $x_{2T} = 0$ is P . Thus, assuming $F > 200/T$, $x_{2T} = 0$ is optimal if $x_{1T} = 1$. Now consider the other case: $x_{1T} = 0$. Now $x_{2T} = 1$ yields $P - F$, while $x_{2T} = 0$ yields at least $P + 100/T - F$. Again, $x_{2T} = 0$ is optimal.

Anticipating this, player one's optimal choice has $x_{1T} = 0$, and the backward induction outcome (5) is established in the usual manner.

3. EXPERIMENTAL DESIGN

The experiment consisted of two sets of five sessions conducted at Penn State University in Summer/Fall 2000. One set employed the AM mechanism with $T = 10$ and $F = 50$, while the other set employed the GP mechanism with the same parameters.

Subjects were respondents to e-mail invitations sent to upper division undergraduates enrolled in the College of Business. The e-mail stated that

sessions would last up to one hour, and participants would receive a participation fee of \$5, plus an opportunity to make additional earnings dependent on decisions. Subjects were not allowed to take part in more than one session.

When subjects arrived for the experiment they were randomly allocated to a computer terminal, and given a set of written instructions, which a monitor then read aloud.⁵ After hearing the instructions subjects were given a quiz. When all subjects had completed the quiz correctly the decision-making part of the experiment began. The only permitted communication between subjects was via formal decisions transmitted by pressing keys on a computer terminal.

In seven of the sessions, eighteen subjects participated in 17 rounds, where each round consisting of a single play of the mechanism game against an anonymous opponent. Subjects were rematched between rounds so that each subject played every other subject once. Two AM sessions and one GP session were conducted with 16 subjects (due to no-shows). These subjects also participated in 17 rounds, playing 13 other subjects once and 2 other subjects twice.⁶ For the GP sessions each subject played nine times in one role and eight in the other, according to a predetermined sequence.

Allowing subjects to repeat a decision task gives them opportunities to learn about the strategic structure of the game. However, because we were primarily interested in the (one-shot) games described in the previous section, we had subjects play different opponents in different rounds in order to reduce repeated game effects.⁷ For example, if the same two subjects were paired in all 17 rounds of the AM mechanism game, then, except for the final round, it would not be true that sequences beginning with β are dominated.

At the beginning of around each subject was informed of the current round number and point earnings, but were not given any information as to the identity of his or her opponent. Subjects then made their decisions according to the rules of the relevant mechanism. At the end of the round each subject was shown the decisions made in the game he or she had played, and his or her own point earnings from the game.

The first three rounds were designated practice rounds: point earnings from these did not count toward final point earnings. Final point earnings were the accumulated point earnings from rounds 4 to 17. These were

⁵ Copies of the instructions are available at <http://www.nyu.edu/jet>.

⁶ In these sessions the experimenter announced that subjects would play 2 other participants twice and the other 13 participants once.

⁷ Clark and Sefton [5] find significant differences between the outcomes of a repeated AM mechanism game and a sequence of one-shot AM mechanism games when the mechanism is designed to implement the payoff-dominant outcome of a stag-hunt game. (The mechanism was more successful in the repeated game environment.)

converted into cash using an exchange rate of 1 cent per point at the end of the session, which was added to the \$5 participation fee. Earnings averaged \$16.97, and ranged from \$3 to \$23.70. Sessions lasted less than one hour, averaging approximately 45 minutes.

4. RESULTS

Iterative dominance and backward induction make sharp predictions in the games that we study. We begin this section by comparing the actual outcomes with these predictions. Fig. 2 graphs the proportion of games resulting in the predicted outcome for each round.

In round 4, the first round that counted toward earnings, not one out of 43 AM mechanism games resulted in the outcome surviving iterative dominance, while only 3 of 44 GP mechanism games ($\approx 7\%$) resulted in the backward induction outcome. By the final round rather more games resulted in the predicted outcome: 30% of the AM mechanism games and 18% of the GP mechanism games. (Note that the AM versus GP ranking based on play in the final round reverses that based on initial play.) Nevertheless, a notable point is that the majority of games do not result in the predicted outcome. Across all rounds for which subjects were paid, 15%




FIG. 2. Proportion of games resulting in predicted outcome.

of AM mechanism games and 10% of GP mechanism games resulted in the predicted outcomes.

This finding should not be interpreted necessarily as evidence of irrational behavior. Consider, for example, the extent to which subjects play dominated strategies. In the AM mechanism game, any strategy with $x_{11} = 1$, i.e., placing β in the first element of the sequence, is dominated. Any other strategy is undominated. In the GP mechanism game any second-mover strategy that results in $x_{2T} = 1$ (i.e., the second mover placing β in the last element of the sequence) is dominated by another strategy that is identical except it plays $x_{2T} = 0$ after any history. Thus we associate any observed second-mover sequence ending with a choice of β with the use of a dominated strategy, and any other second-mover sequence with the use of an undominated strategy. Figure 3 shows the proportion of dominated strategies observed across rounds.

It is clear from Fig. 3 that most observed strategies were undominated, and that the use of dominated strategies declined over the course of the sessions. In fact, across all rounds for which subjects were paid, 95% of submitted AM strategies were undominated, while 98% of (second-mover) GP strategies were undominated. Again the AM versus GP ranking depends upon whether it is based on final or initial play.




FIG. 3. Proportion of dominated strategies.

Taken together, Figs. 2 and 3 show that while dominated strategies are rarely used (especially in later rounds), subjects fail to drive the logic of iterated dominance/backward induction to its conclusion. That subjects use a limited number of iterations of dominance in the AM mechanism games and a limited number of steps of backward induction in the GP mechanism games is borne out by another strong pattern in our data. Note that in the AM mechanism, a subject has 2^{10} pure strategies from which to choose. In the GP mechanism a player's pure strategies can generate 2^{10} distinct sequences. However, for each mechanism only 11 of these possible sequences have the following monotonicity property. A monotonic sequence in the AM mechanism game has, for $k > m$, $x_{ik} \geq x_{im}$, while a monotonic sequence in the GP mechanism game has, for $k > m$, $x_{ik} \leq x_{im}$. For the AM mechanism 92% of the sequences observed in rounds 4–17 are monotonic, while the corresponding percentage for the GP mechanism is 90%. Moreover, just as the frequency of dominated strategies declines over the course of a session, so too does the frequency of non-monotonic sequences (see Fig. 4), particularly in the case of the AM mechanism.

Thus, subjects in the AM sessions typically submitted sequences consisting of a number of α choices followed by a number of β choices, while subjects in the GP sessions tended to do the reverse. In Fig. 5 we show the




FIG. 4. Proportion of nonmonotonic sequences.




FIG. 5. Average number of α choices in monotonic sequences by round.

average number of α choices in a sequence, restricting attention to monotonic sequences. This figure indicates that subjects using monotonic sequences initially choose sequences consisting of, on average, three or four α 's. As the session progresses, they gradually choose more α s so that, by the final round, sequences involve mainly α choices.

Figures 2–5 all exhibit dynamics that move data in the direction of the predicted outcome. These dynamics are consistent with an explanation whereby subjects learn to eliminate more and more conditionally dominated strategies as they repeat the game against different opponents. Even though subjects face different opponents in each round, such a learning process could result from a model whereby subjects tend to change their behavior in the direction of better replies to past opponents' choices, as in Selten and Stoecker's [17] analysis of finitely repeated prisoner's dilemma supergames. If anything, the figures suggest that the rate of learning is faster in the simultaneous-move game.

We close this section with a statistical comparison of the two mechanisms. Both predict the same outcome—albeit using different solution concepts—and both are applied to the same game. For formal statistical tests of mechanism equivalence we first take an individual session as our unit of observation. Thus, for any particular measure of performance, our design yields five observations on each mechanism, which we compare

using a two-tailed Wilcoxon test. We consider several measures of performance. We consider the percentage of games resulting in the predicted outcome, first pooling over all rounds for which subjects were paid (rounds 4–17), and second focussing on the last half of these rounds. We also consider the percentage of α choices in rounds 4 to 17, as well as the percentage in rounds 11–17. The data are presented in Table I.

Averaging over all sessions the AM mechanism outperforms the GP mechanism according to all of these measures. However, the observed difference is only significant at the 10% level in the case of the percentage of α choices in the last seven rounds.

Comparing mechanisms at a more disaggregated level is problematic since individual decisions within a session cannot be taken to be independent. Our design does, however, allow a clean comparison of first-round data, since each first-round game can be considered as independent. The proportions of first-round games resulting in the predicted outcome, 0/43 for the AM and 4/44 for the GP mechanism, are not significantly different at conventional levels (the p -value for Fisher's exact test is 0.116). Likewise, while α choices are more frequently observed in the first-round games of the AM sessions (48%) than the GP sessions (42%), the difference is not significant (p -value for the Wilcoxon test is 0.1183).

Thus, in our laboratory setting the GP mechanism does not outperform its static counterpart. In most cases observed differences between the two mechanisms favor the AM mechanism, although these differences are

TABLE I
Comparison of Performances of Mechanisms

| Session | % of games resulting in predicted outcome (rounds 4 to 17) | % of α choices (rounds 4 to 17) | % of games resulting in predicted outcome (last seven rounds) | % of α choices (last seven rounds) |
|--------------------|--|--|---|---|
| AM1 | 0.0 | 48.5 | 0.0 | 55.9 |
| AM2 | 22.2 | 73.3 | 33.3 | 83.6 |
| AM3 | 12.5 | 72.2 | 23.2 | 78.1 |
| AM4 | 1.8 | 58.5 | 0.0 | 62.0 |
| AM5 | 37.3 | 85.5 | 58.7 | 94.8 |
| GP1 | 11.1 | 64.0 | 12.7 | 69.1 |
| GP2 | 0.0 | 37.1 | 0.0 | 41.6 |
| GP3 | 2.4 | 47.4 | 3.2 | 51.1 |
| GP4 | 6.3 | 48.8 | 9.5 | 53.2 |
| GP5 | 31.0 | 69.8 | 33.3 | 76.1 |
| Wilcoxon p value | 0.6761 | 0.1172 | 0.6761 | 0.0758 |

usually statistically insignificant. The one exception is the percentage of α choices in later rounds: using this measure the AM mechanism has a statistically significant advantage.

5. CONCLUSION

The research goal of identifying simple and practical mechanisms for implementing desired outcomes requires a combination of innovative theory and empirical scrutiny. This paper has reported laboratory experiments to compare two recently proposed mechanisms that use alternative solution concepts to implement desired outcomes. The first mechanism, due to Abreu and Matsushima [1], uses iterative strict dominance, a solution concept that has a strong decision-theoretic foundation. However, when the predicted outcome calls for many iterations of dominance, some commentators have expressed reservations about this concept, and laboratory evidence finds substantial discrepancies between predicted and actual outcomes. The second mechanism, due to Glazer and Perry [7], uses backward induction. This solution concept is, arguably, more intuitive and simpler to understand. However, this concept has also been criticized when many steps of backward induction are required, and it has been found wanting in some laboratory settings. Our experiment applies the two mechanisms to the same implementation task, making possible a direct comparison of their performance.

This task clearly presents a challenge to both mechanisms: the majority of outcomes fail to conform to that predicted by iterative strict dominance or backward induction. However, both mechanisms have promising features. The pervasive use of undominated strategies indicates that subjects are aware of the incentives to avoid being fined, and there is a tendency for subjects to apply more steps of dominance/induction with repetition.

Overall we find that, according to simple measures of performance, the AM mechanism outperforms the GP mechanism, although differences are usually statistically insignificant. Thus, our experiment fails to support the argument that sequential mechanisms will be more successful in implementing desired outcomes.

REFERENCES

1. D. Abreu and H. Matsushima, Virtual implementation in iteratively undominated strategies: complete information, *Econometrica* **60** (1992), 993–1008.
2. D. Abreu and H. Matsushima, A response to Glazer and Rosenthal, *Econometrica* **60** (1992), 1439–1442.

3. R. J. Aumann, Backward induction and common knowledge of rationality, *Games Econ. Behav.* **8** (1995), 6–19.
4. E. Ben-Porath, Rationality, Nash equilibrium and backwards induction in perfect information games, *Rev. Econ. Stud.* **64** (1997), 23–46.
5. K. Clark and M. Sefton, Repetition and signaling: Experimental evidence from games with efficient equilibria, *Econ. Lett.* **70** (2001), 357–362.
6. M. Costa-Gomes, V. P. Crawford, and B. Broseta, Cognition and behavior in normal-form games: An experimental study, *Econometrica*, in press.
7. J. Glazer and M. Perry, Virtual implementation in backwards induction, *Games Econ. Behav.* **15** (1996), 27–32.
8. J. Glazer and R. W. Rosenthal, A note on Abreu–Matsushima mechanisms, *Econometrica* **60** (1992), 1435–1438.
9. J. Glazer and A. Rubinstein, An extensive game as a guide for solving a normal game, *J. Econ. Theory* **70** (1996), 32–42.
10. G. W. Harrison and K. A. McCabe, Testing non-cooperative bargaining theory in experiments, in “Research in Experimental Economics” (R. M. Isaac, Ed.), Vol. 5, pp. 137–169, JAI Press, Greenwich, CT, 1992.
11. R. D. McKelvey and T. R. Palfrey, An experimental study of the centipede game, *Econometrica* **60** (1992), 803–836.
12. R. Nagel, A survey of experimental guessing games: a study of bounded rationality and learning, in “Games and Human Behavior: Essays in Honor of Anmon Rapoport” (D. Budescu, I. Erev, and R. Zwick, Eds.), pp. 105–142, Lawrence Erlbaum Associates, New Jersey, 1999.
13. R. Rosenthal, Games of perfect information, predatory pricing and the chain store paradox, *J. Econ. Theory* **25** (1981), 92–100.
14. A. E. Roth, Bargaining experiments, in “The Handbook of Experimental Economics” (J. H. Kagel and A. E. Roth, Eds.), pp. 253–348, Princeton Univ. Press, Princeton, NJ, 1995.
15. M. Sefton and A. Yavas, Abreu–Matsushima mechanisms: experimental evidence, *Games Econ. Behav.* **16** (1996), 280–302.
16. R. Selten, The chain store paradox, *Theory and Decision* **9** (1978), 127–159.
17. R. Selten and R. Stoecker, End behavior in sequences of finite prisoner’s dilemma supergames, *J. Econom. Behav. Org.* **7** (1986), 47–70.