WHAT DOES IT TAKE TO ELIMINATE THE USE OF A STRATEGY STRICTLY DOMINATED BY A MIXTURE?

John Van Huyck, Frederick Rankin, and Raymond Battalio

Draft mx1_1_4
June 1998

Comments to john.vanhuyck@tamu.edu
Related research available at
http://econlab10.tamu.edu/JVH_gtee

Abstract: This paper reports an experiment to determine whether subjects will learn to stop using a strictly dominated strategy that can be an above average reply. It is difficult to find an experimental design that eliminates the play of the strictly dominated strategy completely. The least effective treatment used money to motivate behavior directly. The most effective treatment uses a binary-lottery to induce preferences, but even this treatment requires giving subjects plenty of experience. Doing so reduced the play of the strictly dominated strategy to around 10 percent by the end of a session. There is no evidence for the explosive cycling needed to make the strictly dominated strategy an above average reply. The most surprising result was observing cycles in the frequency with which the strictly dominated action was played or spite cycles. Spite cycles appear using both evolutionary and fixed pairs protocols.

Key Words: game theory, mixed strategies, dominance, induced value theory, risk aversion, binary lotteries, uncertainty aversion, spite, human behavior.

JEL Classification: c72, c78, c92, d83

Acknowledgments: We thank John Wildenthal for research assistance. Eric Battalio implemented the experimental design on the TAMU economics laboratory network. The National Science Foundation and the Texas Advanced Research Program provided financial support. A faculty development leave at the University of Pittsburgh provided Van Huyck with time to draft this version of the paper.

I. INTRODUCTION

A strictly dominated strategy does worse than the strategy that dominates it in all possible outcomes. It is never in the set of best replies to any feasible belief. Since it can never be a best reply, rational choice requires that a player avoid the use of strictly dominated strategies. Advising players to avoid using strictly dominated strategies does not require any theory of mutually consistent behavior. Eliminating strictly dominated strategies ought to be intuitively appealing to anyone who understands the definition.

The advice to avoid using strictly dominated strategies is sometimes thought to be unsound, because it can conflict with efficiency. Binmore (1994, chapter 3) provides a critical survey of the “circle-squarers” who think it is bad advice. Much of this confusion comes from a failure to distinguish strict dominance arguments from iterated dominance arguments, which require a model of how others behave.

The large experimental literature on repeated dominance solvable games, like the prisoners’ dilemma and many oligopoly and public goods experiments, seems particularly prone to this confusion. The use of evolutionary protocols to implement strategic form games, which allows subjects to learn while mitigating repeated game effects, has proven to be an effective way to observe behavior that is much closer to the predictions of strict dominance. For example, Cooper *et al.* (1995) report that using an evolutionary protocol reduces the play of a strictly dominated strategy in the prisoners’ dilemma game to 12 percent in a fairly short experiment. Of course, observing the use of strictly dominated strategies undermines our confidence in the usefulness of rational choice models for prediction.

The prisoners’ dilemma has been widely studied because it poses a conflict between rationality and efficiency. In this paper, we report an experiment that does not involve a conflict between rationality and efficiency. Rationality and efficiency will both require the elimination of the strictly dominated strategy.

This paper examines the use of a strategy that is not strictly dominated by any pure strategy but is strictly dominated by a mixture of pure strategies. Some players may find the strictly dominated strategy attractive because it is a second best reply to all of the pure strategies in the support of the dominating mixed strategy. It is, however, the worst reply to itself.

General theories of adaptive behavior, like Milgrom and Roberts (1991), predict that learning will eliminate the use of a strategy strictly dominated by a mixture under fairly weak conditions. After eliminating the dominated strategy in the game studied in the experiment what remains is the rock-scissors-paper game. It takes much stronger assumptions to predict that learning will converge to a mixed strategy equilibrium. In the rock-scissors-paper game, some dynamics, especially deterministic dynamics, can get caught in cycles and never converge to the unique equilibrium. For the game studied in the experiment, Dekel and Scotchmer (1992) have proved that for almost all initial conditions, the strictly dominated strategy is not eliminated by a discrete time replicator dynamic. The explosive cycle in the rock-scissors-paper component of the game makes the dominated strategy an above average reply sufficiently often that it never goes extinct.[1]

This paper reports an experiment to determine whether subjects will learn to stop using a strictly dominated strategy that can be an above average reply. It is difficult to find an experimental design that eliminates the play of the strictly dominated strategy completely. The least effective treatments used money to motivate behavior directly. The most effective treatment uses a binary-lottery to induce preferences, but even this treatment requires giving subjects plenty of experience. Doing so reduced the play of the strictly dominated strategy to around 10 percent by the end of a session.

There is no evidence for the explosive cycling needed to make the strictly dominated strategy an above average reply. The most surprising result was observing cycles in the frequency with which the strictly dominated action was played or spite cycles. Spite cycles appear using both evolutionary and fixed pairs protocols.

---

[1] This striking result turns out to depend on the particular way Dekel and Scotchmer discretize the replicator dynamic, see Cabrales and Sobel (1992).

II. ANALYTICAL FRAMEWORK

To focus the analysis, consider the following 2×2 game: *MD*. The game is symmetric and each player has a dominant strategy: *M*. Game *MD* has a dominance solvable equilibrium, *MM*. If the numbers in the table denote the players's von Neumann/Morgenstern utility, then the outcome of the game will be the strategy combination *MM* for any subjective belief the players may hold. Neither player has to construct any model of his opponent's behavior what so ever to conclude he ought to play his part of the equilibrium. Doing so also happens to correspond to selecting the symmetric efficient outcome.

	M	*D*
M	30 30	25 5
D	5 25	0 0

Game *MD*

Rapoport, *et al.* (1976, table 8.1) report that 6 of 10 pairs playing an ordinal variant of game *MD* play strategy combination *MM* in every period for 300 periods. Only 1 of 10 pairs plays *MM* less than 90 percent of the time. The overall average frequency of play for all 10 pairs and 300 periods is 92 percent. We can be confident that the vast majority of subjects play *MM*.

The von Neumann/Morgenstern representation of a player's preferences reduces compound lotteries using probability calculus. Let strategy *M* denote the mixture {1/3,1/3,1/3,0} in an expanded game *RSPD*. The resulting compound lottery is strategically equivalent to the outcome *MM* in game *MD*. Since *D* is dominated by *M* there are no beliefs that can rationalize playing *D* when the numbers in the bi-matrix denote utility.

	R	*S*	*P*	*D*
R	15 15	0 75	75 0	25 5
S	75 0	15 15	0 75	25 5
P	0 75	75 0	15 15	25 5
D	5 25	5 25	5 25	0 0

Game *RSPD*

The game that results after deleting the dominated strategy *D* is the widely discussed rock-scissors-paper (*RSP*) game. The *RSP* game has no pure strategy equilibria and illustrates the possibility of best reply cycles. *R* is a best reply to *S*; *S* is a best reply to *P*; and *P* is a best reply to *R*.[2] The unique equilibrium is in mixed strategies and is the strategy combination *MM*.

A. Adaptive Learning

Adaptive learning may cause players to stop using a strictly dominated strategy, like *D*. Milgrom and Roberts (1991) provide a general theory of adaptive learning, which includes Cournot dynamics, fictitious play, stochastic learning processes based on learning to optimize, like genetic algorithms, and stochastic learning processes based on experimentation, like reinforcement learning. They prove that if behavior is consistent with adaptive learning then it will converge to the set of iteratively undominated

---

[2] Rock breaks scissors. Scissors cut paper. Paper covers rock.

strategies.[3]

It is more difficult to construct models of adaptive learning that not only eliminate the play of dominated strategies, but also converge to the unique mixed equilibrium of the reduced game. Fudenberg and Kreps (1993) and Jordan (1993) are examples. Crawford (1993) summarizes this work by writing, "a fully satisfactory adaptive justification for equilibria in which individual players play mixed strategies is unlikely to be forthcoming."

B. Replicator Dynamics

The continuous time replicator dynamic has proved to be a useful selection dynamic in many situations, because it makes accurate predictions about the influence of matching protocols on the stability of symmetric Nash equilibria and about the relationship between initial conditions and terminal outcomes, see Crawford (1991), Friedman (1991, 1996), Van Huyck, *et al.* (1995), and Van Huyck, Battalio, Rankin (1997). The protocol describing how players interact plays an important role in determining the dimension of the resulting dynamical system and the stability of the dynamical system's fixed points. In an evolutionary game, the constituent game is played repeatedly by n players randomly drawn from populations within a community C.

A strictly dominated strategy becomes extinct under the continuous time replicator dynamic, see Weibull (1995, proposition 3.1). Dekel and Scotchmer (1992) used a game similar to *RSPD* as a counter example to the general proposition that the discrete replicator dynamic eliminates strictly dominated strategies, like D. We briefly review the example, which illustrates the possibility that, if what is best cycles and if the strategy is above average often enough, a strategy that is never a best response may nevertheless survive.[4]

Let the unit coordinate vectors $e_1 \equiv (1,0,0,0)$, $e_2 \equiv (0,1,0,0)$, $e_3 \equiv (0,0,1,0)$, and $e_4 \equiv (0,0,0,1)$ denote the four pure strategies, $R, S, P,$ and D in game *RSPD*. Assuming a single large population with random pairwise

---

[3] A sequence of strategies is "consistent with adaptive learning if player n eventually chooses only strategies that are nearly best responses to some probability distribution over his competitors' joint strategies, where near zero probability is assigned to strategies that have not been played for a sufficiently long time. (p.85)"

[4] Hofbauer and Sigmund (1988) and Weissing (1991) analyze cycling in the rock-scissors-paper game.

matching, let s_i^t denote the fraction of the population using strategy *i* at time *t* and let s^t denote the vector $(s_1^t, s_2^t, s_3^t, s_4^t)$. All feasible population frequency vectors s^t lie in the simplex $\mathbf{S}^4$.

The expected payoff to a player using strategy *i* is $e_i.R.s^t$. The average expected payoff in state s^t is $s^t.R.s^t$. The discrete replicator dynamics for payoff matrix *R* with one population is given by the following system of difference equations:

$$\frac{s_i^{t+1}}{s_i^t} = \frac{e_i.R.s^t}{s^t.R.s^t} \qquad \forall\, i.$$

For game *RSPD*, a fixed point of the replicator dynamics is the point corresponding to the unique equilibrium, (1/3,1/3,1/3,0). Dekel and Scotchmer (1992) prove that this fixed point is asymptotically unstable. Specifically, they prove that for a solution path s^t starting from a completely mixed initial condition, s_4^t converges to zero if and only if $s_1^0 = s_2^0 = s_3^0$.

While strategy *D* is dominated by *M*, it is an above average reply when ever the state is close to one of the other three pure strategies. Figure 2 graphs the wire frame of three polyhedra, one for *R*, *S*, *P* respectively, for those states in which *D* is an above average reply, that is,

$$e_4.RSPD.s^t > s^t.RSPD.s^t$$

under the restriction that $s_4 = 1 - s_1 - s_2 - s_3$. The wire frames describe one above average reply polyhedra for each of three vertices associated with *R*, *S*, and *P* respectively. The frequency of strategy *D* is increasing for all states within the above average polyhedra. The interior of the space between the mixed equilibrium, denoted with a dot, and the above average pohlyhedra describe states in which *D* is a below average reply and its frequency is decreasing in the population.

<insert figure 2 about here>

Figure 2 graphs an example nonconvergent solution path in $\mathbf{S}^4$. The path begins close to *D*. Initially, the frequency of *D* declines, but once the state is sufficiently close to the face of the simplex in which *D* is extinct the path begins to cycle explosively. This explosive cycling moves the state into one of the polyhedra in which *D* is an above average reply. While the state is within an above average polyhedra the frequency of *D* grows. Eventually the state passes out of one polyhedra and moves to the next as the cycling amongst *R, S,* and *P* continues. During these transitions the frequency of *D* declines, but not by enough to prevent it from staying between 45 and 58 percent of the strategies in the population.

III. EXPERIMENTAL DESIGN

Table one summarizes the experimental design. Human subjects played game *RSPD* either as an evolutionary game, cohorts 1 to 16, or as a repeated game, cohorts 17 and 18. Six subjects participated in each cohort using a one population evolutionary matching protocol and twelve subjects participated in each cohort using a repeated pairs matching protocol.

The subjects had complete information about both their own and everybody else's payoff matrix. Cohorts 1 to 5 used a matrix presentation of the payoffs, which focuses attention on the dominance relation and de-emphasizes mutual consistency. Subjects were instructed on how to use the symmetry of the problem to deduce the payoffs of the other participant. Cohorts 6 to 18 use a bi-matrix presentation of the payoffs, which focuses attention on mutual consistency.

Preferences were induced using either money payoffs directly, all cohorts except 11 to 16, or using a binary lottery, cohorts 11 to 16.[5] Cell entries either denoted cents or the probability of winning a dollar. Probability was explained as the number of chances of winning a dollar out of 100.

In cohorts using the evolutionary matching protocol, the subjects's actions were randomly matched each period to determine an outcome. The subjects were informed that they were being randomly paired. Since outcomes were reported privately, subjects could not use common information about the outcomes in previous periods to coordinate behavior. Subjects confronted an anonymous participant each period in the evolutionary protocol.

In the repeated pairs cohorts, 17 and 18, subjects were randomly matched at the start of the first period and remained matched with the same participant for the duration of the session. The repeated pairs cohorts are included for comparison to the literature. Many learning theories assume

---

[5] See Roth and Malouf (1979) on the use of binary lotteries to induce risk preferences. O'Neil (1987) emphasizes the use of two outcome protocols when investigating mixed strategy equilibria, see also Brown and Rosenthal (1990), Rapoport and Boebel (1992), Shachat (1996), Wooders and Shachat (1997), and Binmore *et al.* (1997).

repeated pairs and most experiments investigating mixed strategies use repeated pairs. Moreover, we conjectured that, unlike matching pennies or rock-scissors-paper, game *RSPD* allows a role for "teaching" the other participant to avoid using *D*. Notice that "teaching" involves the use of a repeated game strategy and that adaptive learning theories are usually defined over stage game rather than repeated game strategies.

To summarize, the experiment used four treatments to study behavior in game *RSPD*: a matrix/cents treatment, a bi-matrix/cents treatment, a bi-matrix/ probability treatment, and a repeated pairs treatment.

The period game was described to the subjects using a computer assisted graphical user interface available in the TAMU economics laboratory. The instructions were read aloud to insure that the description of the game was common information. After reading the instructions, but before the session began, the subjects filled out a questionnaire to determine that they understood how to compute payoffs for themselves and their opponents. If any subject made a mistake on the questionnaire, the relevant section of the instructions was read again.

The subjects were recruited from undergraduate classes in economics and business at Texas A&M University. A total of 120 subjects participated in 18 cohorts. The sessions took about one and one-half hours to conduct. Repeated play of the unique stage game equilibrium would have resulted in each subject earning $27.

IV. EXPERIMENTAL RESULTS

The experimental results are reported in four sections. Aggregate results are reported in section IV.*A*. Treatment effects are reported in section IV.*B*. Cohort effects are examined in section IV.*C*. Individual subject results are reported in section IV.*D*.

A. Aggregate Results

Of the 10,800 choices made by the 120 subjects, 2,163 or about 20 percent were the strictly dominated strategy *D*. The remaining choices were as follows: strategy *R* was chosen 2,843 times, strategy *S* was chosen 3,041 times, and strategy *P* was chosen 2,753 times. The null hypothesis of equilibrium behavior implies that the mean frequency of these three strategies would have been 3,600. The $\chi^2(3)$ statistic for equilibrium play is 145.4, which rejects at all conventional levels of statistical significance.

A null hypothesis of random behavior does slightly better with a $\chi^2(3)$ statistic of 9.9, but this is large enough to also reject random behavior at the 5 percent level of statistical significance. The maximum likelihood estimate of the logit equilibrium gives a noise parameter of 0.06. Given this noise parameter, the logit equilibrium is {0.267,0.267,0.267,0.198}, which gives a $\chi^2(3)$ statistic of 1.02 with a probability value of 0.41. The test fails to reject the logit equilibrium after it has been fitted to the aggregate data.[6] As we shall see aggregating over treatments, cohorts, subjects, and time hides important variations in the data.

B. Treatment Effects

The matrix/cents treatment focuses subjects on the fact that *M* dominates *D* at least in the sense of expected earnings. In the first period of the matrix/cents treatment, 9 of 30 subjects or 30 percent chose *D*. Choosing *D* was the modal action, with each of the other three actions chosen by 7 subjects. Over the first ten periods, action *D* accounts for 92 of 300 actions or 31 percent and is the modal choice, see table two. The other three actions were chosen with about equal frequency.

The subjects attraction to *D* presents a puzzle. One reaction might be to conclude that subjects play dominated strategies and leave it at that. This does not seem to be a terribly constructive reaction. Instead, we considered the following hypotheses:

H1: The subjects attraction to *D* could be due to confusion. If so, we would expect to see the behavior become extinct with experience.

H2: The pairs treatment allows subjects to teach the subject with whom they are matched to avoid *D*. This may reduce confusion and increase the subjects understanding of the usefulness of finding a mutually consistent way to play.

H3: The attraction of *D* might be due to the subjects not imposing a mutual consistency condition on their beliefs. Action *D* is the second best response against three of four pure strategies and a naive analysis might suggest it being satisfactory most of the time. If so, we would expect the bi-matrix presentation, which focuses subjects attention on the other player, to reduce the frequency of *D*.

H4: The attraction of *D* might be due to risk aversion. If so, we would expect the binary lottery to reduce the frequency of *D*, because it

---

[6] See McKelvey and Palfrey (1995) on the logit equilibrium. Briefly, the logit equilibrium is a fixed point of the logit response function: $\text{prob}[e_i,\lambda,q] = \text{Exp}[\lambda\, e_i.RSPD.q] / \Sigma_j \text{Exp}[\lambda\, e_j.RSPD.q]$, where λ is the noise parameter. When λ is zero the logit equilibrium is the centroid of the simplex and as λ goes to infinity a logit equilibrium goes to a Nash equilibrium.

controls for risk aversion.

Experience does reduce play of *D*. For the matrix/cents treatment, play of *D* falls from about 31 percent in the first ten periods to 24 percent in the last ten periods, compare tables two and three. The $\chi^2(1)$ statistic for the null hypothesis of no change in *D* frequency across the two time periods is 3.36, which rejects at the 10 percent level of statistical significance. The probability value is 0.067, see table four. There is weak evidence for a reduction in the frequency of the strictly dominated strategy in the matrix/cents treatment.

The decline in play of *D* is slightly more pronounced in the bi-matrix/cents and bi-matrix/probability treatments. For the bi-matrix/cents treatment, play of *D* falls from 26 percent in the first ten periods to 19 percent in the last ten periods, compare tables two and three. The $\chi^2(1)$ statistic for the null hypothesis of no change in *D* frequency across the two time periods is 4.22, which rejects at the 5 percent level of statistical significance, see table four.

For the bi-matrix/probability treatment, play of *D* falls from 20 percent in the first ten periods to 10 percent in the last ten periods, compare tables two and three. The $\chi^2(1)$ statistic for the null hypothesis of no change in frequency of *D* across the two time periods is 14.12, which rejects at all conventional levels of statistical significance, see table four. The decline in the frequency of *D* is statistically significant at some conventional level for all evolutionary treatments suggesting some merit to H1. However, the size of the decline is, at least to us, surprisingly small. Ninety periods of experience only reduces the frequency of *D* by 7 percent in the matrix/cents and bi-matrix/cents treatments and by 10 percent in the bi-matrix/probability treatment.

There is no evidence for the teaching hypothesis, H2. In the first ten periods of the pairs treatment, *D* was chosen 59 times or about 25 percent and in the last ten periods of the pairs treatment, *D* was chosen 57 times or about 24 percent, compare tables two and three. The $\chi^2(1)$ statistic is 0.045 with a probability value of 0.831, see table four. Requiring subjects to construct *M* from *R*, *S*, and *P* triples the *D* frequency from 8 to 24 percent when compared to the results of Rapoport *et al.* (1976).

There is no statistically significant evidence for the mutual consistency hypothesis: H3. In the first ten periods, the difference between *D* being chosen 31 percent of the time in the matrix/cents treatment and 26 percent of the time in the bi-matrix/cents treatment gives a $\chi^2(1)$ statistic of 1.61, see table four below the diagonal. The probability value of 0.205 does not reject the null hypothesis of no treatment effect. In the last ten periods, the frequency of *D* is 24 and 19 percent for the matrix/cents and bi-matrix/cents treatments respectively. The $\chi^2(1)$ statistic of 2.22 is reported in table four above the diagonal. The probability value of 0.136 does not reject the null hypothesis of no treatment effect at conventional significance levels.

There is evidence for the risk aversion hypothesis: H4. All of the comparisons between the bi-matrix/probability treatment and the other evolutionary treatments using cents are statistically significant at some conventional level, see table four. The most dramatic difference is with the matrix/cents treatment in the last ten periods. In the matrix/cents treatment *D* is chosen 24 percent of the time and in the bi-matrix/probability treatment *D* is chosen 10 percent of the time. The $\chi^2(1)$ statistic of 23.43 is statistically significant at all conventional levels of significance.

The economic significance of this 14 percent difference is best evaluated in the payoff space. Table five reports the mean and standard deviation of earnings by treatment. Subjects earn $2.64 and $1.94 more in the bi-matrix/probability treatment than in the matrix/cents treatment and the bi-matrix/cents treatment respectively. These differences are statistically significant at the 1 percent and 5 percent levels respectively.

Subjects earn from $3.86 to $6.51 less than their equilibrium earnings, that is, subjects earn 75 to 86 percent of equilibrium earnings, see table five. These economically significant differences are also statistically significant at all conventional levels. The matrix/cents treatment has the lowest mean earnings. The standard deviation of earnings is largest in the bi-matrix/probability treatment. The range of earnings in the bi-matrix/probability treatment is $20 (from $15 to $35).

C. Cohort Effects: above average cycles or spite cycles?

The play of a strictly dominated strategy may not be surprising to readers familiar with the Dekel/Scotchmer discrete replicator dynamic, because it predicts the survival of such strategies in game *RSPD*. After a sufficiently long time action *D* will be an above average reply most of the time. The Dekel/Scotchmer prediction turns out not to be a satisfying explanation. Figures 3, 4, and 5 graph the moving average play path of cohorts 1, 2, and 3 to illustrate how different the observed play paths are from the predicted discrete replicator dynamic, compare with figure 2. For the time aggregated data, action *D* is never an above average reply. Subjects who are not playing *D* behave sufficiently randomly that the population frequencies of *R*, S, and *P* are close to being equal. The subjects don’t exhibit the kind of correlated mistakes modeled by the dynamic.

In the evolutionary treatments without time aggregation, *D* is an above average reply to the cohort's state only 7.5 percent of the time. These states occur randomly, that is, the predicted explosive cycles over *R*, *S*, *P* don't materialize. The attraction of *D* can not be due to it being an above average reply, since behavior doesn't cycle through *R*, *S*, and *P*.

A remarkable feature of the data is the large number of cohorts with significant fluctuations in *D* frequency. Figures 6, 7, and 8 graph the 10 period *D* frequency moving average by cohort for the matrix/cents, bi-matrix/cents, and bi-matrix/probability treatments respectively. A typical case is cohort 3. Play of *D* reaches 35 percent in periods 16 to 25 falls to a low of 17 percent in periods 41 to 50 and then rises to 28 percent in periods 80 to 89. A more dramatic case is cohort 10. Play of *D* reaches 41 percent in periods 17 to 28 falls to a low of 17 percent in periods 33 to 42 and then rises to 50 percent in periods 51 to 60 falls to 12 percent in periods 67 to 76 and finally rises back to 28 percent in periods 77 to 86.

Since these cycles can't be explained as above average response cycles, we call them spite cycles. The sense in which these cycles involve spite is that in some cases a subject will respond to a long sequence of his opponent(s) choosing *D* with *D*. Both players earn 0 in the outcome *DD*, but it appears that some subjects prefer an outcome in which both subjects earn 0 to an outcome in which they earn 5 cents and the other participant earns 25 cents, which is what happens against *D* when they choose an undominated strategy.

In the pairs data this phenomena can be rationalized as "teaching" one's opponent to stop playing *D*. Subjects respond to *D* with *D* in the fixed pairs protocol twice as often as they do in the evolutionary protocol. Specifically, in the evolutionary protocol subjects respond to *D* by playing *D* in the following period 22 percent of the time while in the fixed pairs protocol this occurs 43 percent of the time. However, teaching your opponent a lesson seems to be a good way to enter a spite cycle.

The pairs data graphed in figures 9 and 10 exhibit the most spectacular spite cycles. Play of *D* in pair 1 of cohort 17 starts at 40 percent in periods 1 to 10 falls to 10 percent in periods 13 to 22 climbs to 85 percent in periods 38 to 47 and then falls to 0 in periods 55 to 89. Pair 3 of cohort 18 also exhibits spectacular cycles. Play of *D* in pair 3 falls to 60 percent in periods 8 to 17 rises to 90 percent in periods 18 to 27 falls to 0 in periods 48 to 70 and then rises to 45 percent in periods 81 to 90. Play of *D* in pair 2 of cohort 18 starts at 10 percent in the first ten periods rises to 60 percent in periods 30 to 39 falls to 15 percent in periods 50 to 59 and rises to 65 percent in periods 81 to 90.

Pair 6 of cohort 18 is a remarkable pair. They managed to avoid spite cycles by coordinating on a correlated equilibrium of the repeated game. They achieved this coordination in three phases: an initial phase of five periods in which the subjects played *D* once or twice, a transition period of 11 periods in which initially one subject played *P* repeatedly and then the other subject chose *S* repeatedly; finally an equilibrium phase from period 16 to 90 in which one player played *S* in every period and the other player played *R* in even periods and *P* in odd periods. So for the last 75 periods each subject was earning on average 37.5 cents per period or 7.5 cents per period more than the equilibrium of the one shot game. No other pair coordinated on a recognizable correlated equilibrium.

Table six reports the results of a Tukey test for cohort effects in the frequency of strategy *D*. The test groups cohorts based on the criterion that the means are not significantly different from each other. There is substantial evidence of cohort effects. These cohort effects are sufficiently large that they can overcome the treatment effects discussed in section IV.*C*. For example, Tukey grouping E includes cohorts from all three of the evolutionary treatments. However, groupings A and B do not include cohorts from the bi-matrix/probability treatment and grouping G does not include cohorts from either of the cents treatments. This is further statistical evidence of a difference in behavior between the cents treatments and the probability treatments.

It is possible to test this hypothesis directly using a Mann-Whitney test statistic.[7] For the null hypothesis of no difference between the matrix/cents and bi-matrix/cents treatment gives an insignificant test statistic at more than the 20 percent level. Hence, we accept the null and pool the two cents treatments. Table six ranks the cohort means from lowest to highest.

For the null hypothesis that the the cents treatment means are less than or equal to the probability treatment means the Mann-Whitney test statistic is 27, where the sample size of the probability treatment is 6 and of the cents treatment is 10. The test statistic of 27 rejects the null hypothesis at less than 0.01 level. The probability treatment means are significantly less than the cents treatment means.

---

[7] We thank John Kagel for suggesting the Mann-Whitney test as an alternative to the chi-square tests used in table four. The Mann-Whitney test has the advantage that it does not suffer from a lack of independent measures as the chi-square test does nor does it require a lack of cohort effects as the Smirnov test does. The results are robust to these objections.

D. Individual Subject Results

Are the aggregate results representative of individual behavior? Perhaps most subjects recognize that *D* is a dominated strategy, while a few subjects make sub-additive probability judgements or some other judgement error that causes a systematic bias in favor of *D*. Table seven reports the distribution of subjects by the frequency with which they played *D*.

Only 3 of 120 subjects never play *D*. If we add subjects who play *D* 10 percent or less, this comes to 40 of 120 or 33 percent who don't play *D* often. Conversely, no subject ever plays *D* exclusively. In fact, only 9 subjects or 8 percent play *D* more than 40 percent of the time. This leaves 71 subjects or 59 percent who play *D* between 11 and 40 percent of the time. Most subjects are mixing in about the same range as reported in the aggregate and cohort results sections.

Fisher exact tests of individual behavior reject the mixed strategy hypothesis for seventy of the ninety-six subjects in the evolutionary game sessions at the five percent level of statistical significance. So, as many as twenty-six subjects may have been using the dominating strategy *M*.

Table seven suggests treatment differences in the cumulative distribution of individual subject's *D* frequency. The Smirnov test can be used to check for statistically significant differences in the cumulative distribution. Rather than work with the binned data summarized in table seven, we compute the Smirnov statistic for the empirical cumulative distribution function. Let $S_m(x)$ denote a treatments observed cumulative distribution of individual subject's *D* frequency, that is, $S_m(x)=k/m$ where m is the number of subjects in the treatment and k is the count of data equal to or less than x. The Smirnov test statistic is $\max[\ |S_m(x) - S_n(x)|\ ]$.

The Smirnov statistic for the null hypothesis that the matrix/cents treatment and the bi-matrix/probability treatment were drawn from the same population distribution function is 0.42. Given samples of size 30 and 36 respectively, this statistic rejects the null hypothesis at the 1 percent level of statistical significance.

The Smirnov statistic for the bi-matrix/cents and bi-matrix/probability treatments is 0.30. Given sample sizes of 30 and 36, this statistic is just barely significant at the 10 percent level.

The Smirnov statistic for the pairs and bi-matrix/probability treatments is 0.40. Given samples of size 24 and 36 respectively, this statistic rejects at the 5 percent level of statistical significance.

The Smirnov statistic fails to detect significant differences between the empirical cumulative distribution functions of individual subject's *D* frequency for any combination of the three treatments using cents.

V. SUMMARY AND CONCLUSION

In all treatments we observe persistent play of a strictly dominated strategy, *D*. This phenomena, while predicted by the discrete replicator dynamics, can not be explained by it, because the explosive cycles through *R, S, P,* which make *D* an above average reply, never occur. *D* is almost never an above average reply.

In other work, we have found the replicator dynamic to be a useful selection dynamic. In this paper, it doesn't make accurate predictions. A lesson from the experiment is that one should discount models that predict deterministic cycles. Humans just don't seem to get caught in them. It follows then that when a simple deterministic dynamic predicts deterministic cycles the situation ought to be re-examined using a probabilistic choice learning dynamic.

Experience reduced the frequency of *D* in all treatments except the pairs treatment. Somewhat paradoxically, the ability to "teach" the other participant to stop using *D* did not reduce the frequency with which *D* was played over time.

Using a binary lottery significantly reduced the frequency of *D*. The emphasis put on the use of binary lotteries when testing mixed strategy equilibria in the literature seems well founded. The *RSPD* game should be added to the examples in Prasnikar (1993) and Reitz (1993) who have previously compared the performance of induced value techniques and found it important to control for unobserved risk preferences using a binary lottery, see Roth (1995, 81-83).[8]

We observed cycles over time in the frequency of *D,* which we called spite cycles. Pairs seemed particularly prone to spite cycles, but the other treatments also exhibit some cycling. The least cycling occurs in the binary lottery treatment. A bad outcome can be attributed to luck rather than necessarily due to the other participants behavior in the binary lottery treatment. A possibility worth considering in future work is that the binary lottery may be effective not because it controls risk preference, but rather because the binary lottery reduces the likelihood of spite cycles.

It is possible to find a noise parameter for the logit equilibrium such that our non-parametric tests fail to reject the hypothesis that the data was generated by the fitted logit equilibrium. We were surprised by the small size

---

[8] John Kagel in personal communication tells us that he has used the binary lottery technique in signaling games and found an increase in the play of strictly dominated strategies.

of the noise parameter, it was only one third as large as the smallest parameter we had previously estimated. Nevertheless, we think the probabilistic choice framework provides a promising way to organize experimental data.

REFERENCES

Ken Binmore, Chris Proulx, and J. Swierzbinski, "Does Minimax Work? an experimental study," laser-script 1997.

J. Brown and R. Rosenthal, "Testing the Minimax Hypothesis: A re-examination of O'Neill's experiment," *Econometrica*, 58 1990, 1065-1081.

Antonio Cabrales and Joel Sobel, "On the Limit Points of Discrete Selection Dynamics," *Journal of Economic Theory*, 57(3) August 1992, 407-419.

Vincent P. Crawford, "An 'Evolutionary' Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games and Economic Behavior*, 3 1991, 25-60.

Vincent P. Crawford, "Introduction" *Games and Economic Behavior*, 5(3), July 1993, 315-19.

Eddie Dekel and Susan Scotchmer, "On the Evolution of Optimizing Behavior," *Journal of Economic Theory* 57 1992.

Daniel Friedman, "Evolutionary Games in Economics," *Econometrica*, 59(3), May 1991, 637-666.

Daniel Friedman, "Equilibrium in Evolutionary Games: Some Experimental Results," *Economic Journal* 106(434), January 1996, 1-25.

Drew Fudenberg and David M. Kreps, "Learning Mixed Equilibria," *Games and Economic Behavior*, 5(3) July 1993, 320-367.

Josef Hofbauer and Karl Sigmund, *The Theory of Evolution and Dynamical Systems: mathematical aspects of selection*, (Cambridge, UK: Cambridge University Press, 1988).

J.S. Jordan, "Three Problems in Learning Mixed-Strategy Nash Equilibria," *Games and Economic Behavior*, 5(3) July 1993, 368-86.

Paul Milgrom and John Roberts, "Adaptive and Sophisticated Learning in Normal Form Games," *Games and Economic Behavior*, 3(1991), 82-

100.

Richard D. McKelvey and Thomas R. Palfrey, “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, 1995.

Barry O’Neill, “Nonmetric Test of the Minimax Theory of Two-person Zerosum Games,” *Proceedings of the National Academy of Sciences* 84 1987, 2106-2109.

Vesna Prasnikar, “Binary Lottery Payoffs: Do they control risk aversion?” laser-script 1993.

Amnon Rapoport and Richard Boebel, “Mixed Strategies in Strictly Competitive Games: A further test of the minimax hypothesis,” *Games and Economic Behavior* 4(2), April 1992, 261-283.

Anatol Rapoport, Melvin Guyer, and David Gordon, *The 2×2 Game*, (Ann Arbor, MI: University of Michigan Press, 1976).

Thomas Rietz, “Implementing and testing risk-preference-induction mechanisms in experimental sealed-bid auctions,” *Journal of Risk and Uncertainty*, 7(2), October 1993, pages 199-213

Alvin E. Roth, “Introduction” in *Handbook of Experimental Economics*, ed. Kagel and Roth, (Princeton, NJ: Princeton University Press, 1995).

Alvin E. Roth and M. Malouf, “Game-theoretic Models and the Role of Information in Bargaining,” *Psychological Review* 1979, 574-594.

Jason Shachat, “Mixed Strategy Play and the Minimax Hypothesis,” laser-script, 1996.

J. Van Huyck, R. Battalio, S. Mathur, A. Ortmann and P. Van Huyck, “On the Origin of Convention: Evidence from Symmetric Bargaining Games,” *International Journal of Game Theory*, 24(2), 1995, 187-212.

John Van Huyck, Raymond Battalio, and Frederick Rankin, “On the Origin of Convention: Evidence From Coordination Games,” with R. Battalio and F. Rankin, *The Economic Journal*, 107(442) May 1997, 576-97.

John Van Huyck, Joseph Cook, and Raymond Battalio, "Adaptive Behavior and Coordination Failure," *Journal of Economic Behavior and Organization*, 32(4), 1997, 483-503.

Jörgen W. Weibull, *Evolutionary Game Theory,* (Cambridge,MA: The MIT Press, 1995)

F. Weissing, "Evolutionary stability and dynamic stability in a class of evolutionary normal form games," in *Games Equilibrium Models. I Evolution and Game Dynamics*, (Berlin/New York: Springer-Verlag, pp. 29-97,1991).

John Wooders and Jason M. Shachat, “On the Irrelevance of Risk Attitudes in Repeated Two-Outcome Games,” laser-script October 1997.

Table One
Experimental Design

Cohort	Matching Protocol	Earnings Table	Units	Cohort Size
1	Evolutionary	Matrix	Cents	6
2	Evolutionary	Matrix	Cents	6
3	Evolutionary	Matrix	Cents	6
4	Evolutionary	Matrix	Cents	6
5	Evolutionary	Matrix	Cents	6
6	Evolutionary	Bi-Matrix	Cents	6
7	Evolutionary	Bi-Matrix	Cents	6
8	Evolutionary	Bi-Matrix	Cents	6
9	Evolutionary	Bi-Matrix	Cents	6
10	Evolutionary	Bi-Matrix	Cents	6
11	Evolutionary	Bi-Matrix	Prob.	6
12	Evolutionary	Bi-Matrix	Prob.	6
13	Evolutionary	Bi-Matrix	Prob.	6
14	Evolutionary	Bi-Matrix	Prob.	6
15	Evolutionary	Bi-Matrix	Prob.	6
16	Evolutionary	Bi-Matrix	Prob.	6
17	Repeated Pairs	Bi-Matrix	Cents	12
18	Repeated Pairs	Bi-Matrix	Cents	12

Table Two

Period 1 to 10 Choice by Treatment Contingency Table

Frequency and Row Percent

$\chi^2(9) = 12.635$ with a p-value of 0.18

Treatment	*R*	*S*	*P*	*D*	Total
Matrix	64	72	72	92	300
Cents	21.33	24.00	24.00	30.67	
Bi-matrix	71	81	70	78	300
Cents	23.67	27.00	23.33	26.00	
Bi-matrix	86	107	95	72	360
Probability	23.89	29.72	26.39	20.00	
Pairs	56	57	68	59	240
Cents	23.33	23.75	28.33	24.58	
Total	277	317	305	301	1200
	23.08	26.42	25.42	25.08	

Table Three

Period 81 to 90 Choice by Treatment Contingency Table

Action Frequency and Row Percent

$\chi^2(9) = 43.37$ with a p-value of 0.001

Treatment	*R*	*S*	*P*	*D*	Total
Matrix	94	75	59	72	300
Cents	31.33	25.00	19.57	24.00	
Bi-matrix	91	75	77	57	300
Cents	30.33	25.00	25.67	19.00	
Bi-matrix	95	100	129	36	360
Probability	26.39	27.78	35.83	10.00	
Pairs	65	63	55	57	240
Cents	27.08	26.25	22.92	23.75	
Total	345	313	320	222	1200
	28.75	26.08	26.67	18.50	

Table Four

$\chi^2(1)$ statistics and probability values for treatment differences in the frequency of *RSP* versus *D*

	Matrix Cents	Bi-matrix Cents	Bi-matrix Probability	Pairs Cents
Matrix Cents	*3.36*	2.22	**23.43**	0.005
	0.067	0.136	0.001	0.946
Bi-matrix Cents	1.61	**4.22**	**10.95**	1.81
	0.205	0.040	0.001	0.179
Bi-matrix Probability	**9.97**	*3.35*	**14.12**	**20.79**
	0.002	0.067	0.001	0.001
Pairs Cents	2.45	0.14	1.77	0.045
	0.118	0.707	0.183	0.831

Notes: Along the diagonal compares first 10 and last 10 periods within a treatment. Across treatment comparisons for the first 10 periods are below the diagonal and for last 10 periods are above the diagonal. $\chi^2(1)$ statistics with a probability value below the 10 percent level are indicated in *italics* and with a probability value below the 5 percent level are indicated in **bold**.

Table Five

Mean Earnings by Treatment

And Deviations from Equilibrium Prediction of $27

Treatment	Mean (-27)	Std (t-statistic)	Minimum	Maximum
Matrix Cents	$20.49	$2.44	$16.00	$25.15
	(-6.51)	(-14.59)		
Bi-matrix Cents	$21.19	$2.83	$15.30	$26.35
	(-5.81)	(-11.23)		
Bi-matrix Probability	$23.14	$4.41	$15.00	$35.00
	(-3.86)	(-5.26)		
Pairs Cents	$21.48	$5.24	$14.00	$32.05
	(-5.52)	(-5.16)		

Table Six
Cohort Mean Frequencies Grouped by
Tukey's Studentized Range Test
(Cohort means with the same letter are not statistically different)

Treatment	Cohort	Rank	Frequency of *D*	Tukey Grouping			
Matrix/Cents	1	16	.2963	A			
Bi-matrix/Cents	10	15	.2796	A	B		
Matrix/Cents	2	14	.2778	A	B		
Matrix/Cents	4	13	.2648	A	B		
Bi-matrix/Cents	9	12	.2593	A	B		C
Matrix/Cents	3	11	.2259	A	B	D	C
Bi-matrix/Cents	6	10	.2130		B	D	C
Bi-matrix/Probability	13	9	.2074	E	B	D	C
Bi-matrix/Probability	11	8	.1870	E	F	D	C
Matrix/Cents	5	7	.1759	E	F	D	
Bi-matrix/Cents	7	6	.1593	E	F	D	
Bi-matrix/Cents	8	5	.1593	E	F	D	
Bi-matrix/Probability	12	4	.1352	E	F		G
Bi-matrix/Probability	15	3	.1130		F		G
Bi-matrix/Probability	16	2	.0815				G
Bi-matrix/Probability	14	1	.0741				G

Table Seven
Distribution of Subjects by Frequency Bins

D Frequency	matrix cents	bi-matrix cents	bi-matrix probability	pairs	total
0		1	2		3
1-10	6	8	17	6	37
11-20	6	8	9	3	26
21-30	11	5	4	8	28
31-40	3	5	3	6	17
41-50	1	1		1	3
51-60	2	1	1		4
61-70		1			1
71-80					
81-90	1				1
91-100					
total	30	30	36	24	120

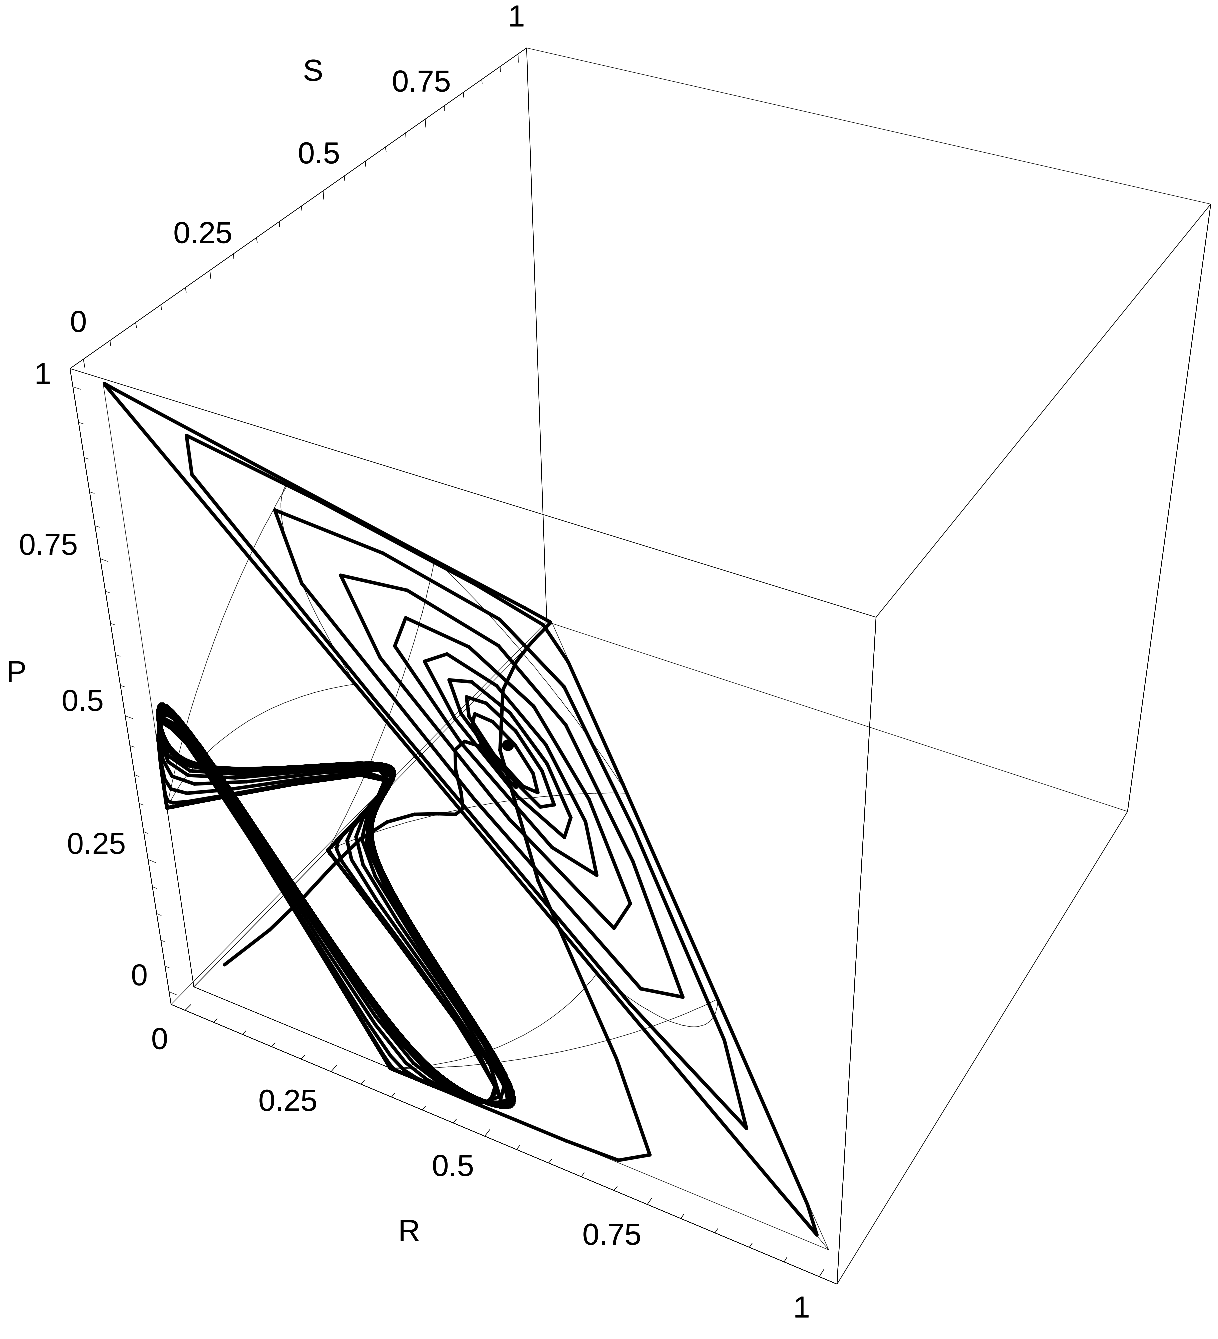

Figure 2: A nonconvergent solution of the discrete replicator dynamic for game *RSPD*. The initial condition is {0.036, 0.03, 0.034, 0.9}. Thin lines represent the wire frame of the simplex and above average payoff polyhedra of *D*. The black dot denotes the unique equilibrium {1/3,1/3,1/3,0}.

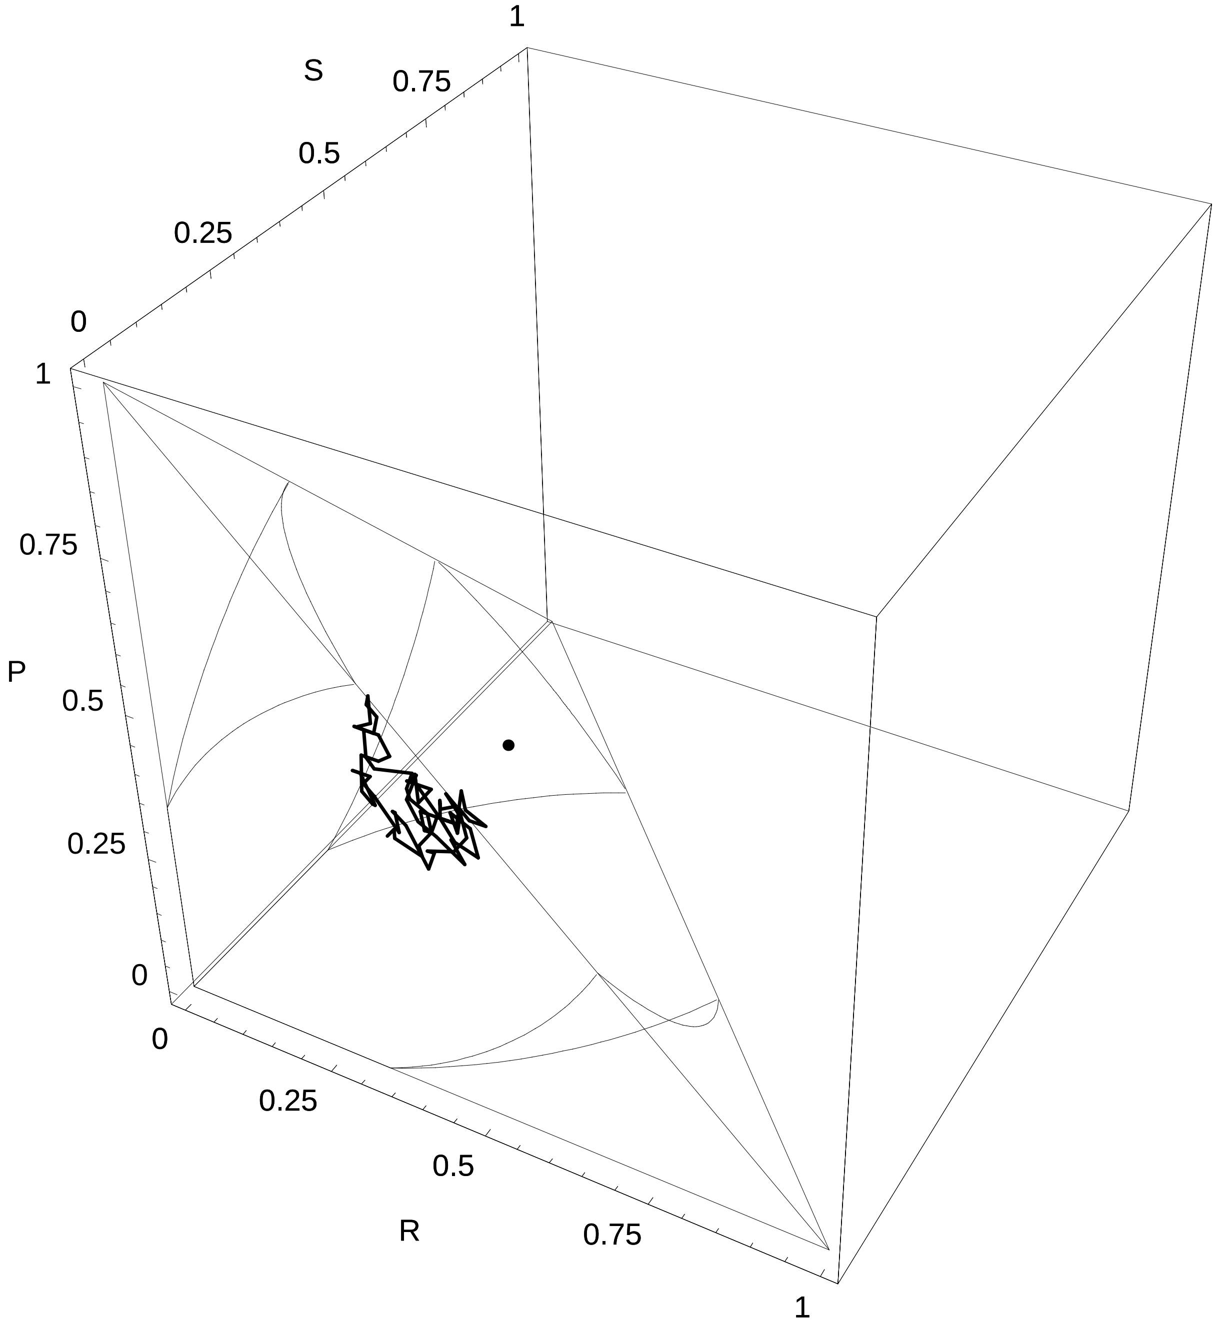

Figure 3: 10 Period Moving Average of Cohort 1 play path plotted in the simplex. Notice that bold line denoting data never penetrates wire frame polyhedra denoting regions of the state space in which *D* is an above average reply.

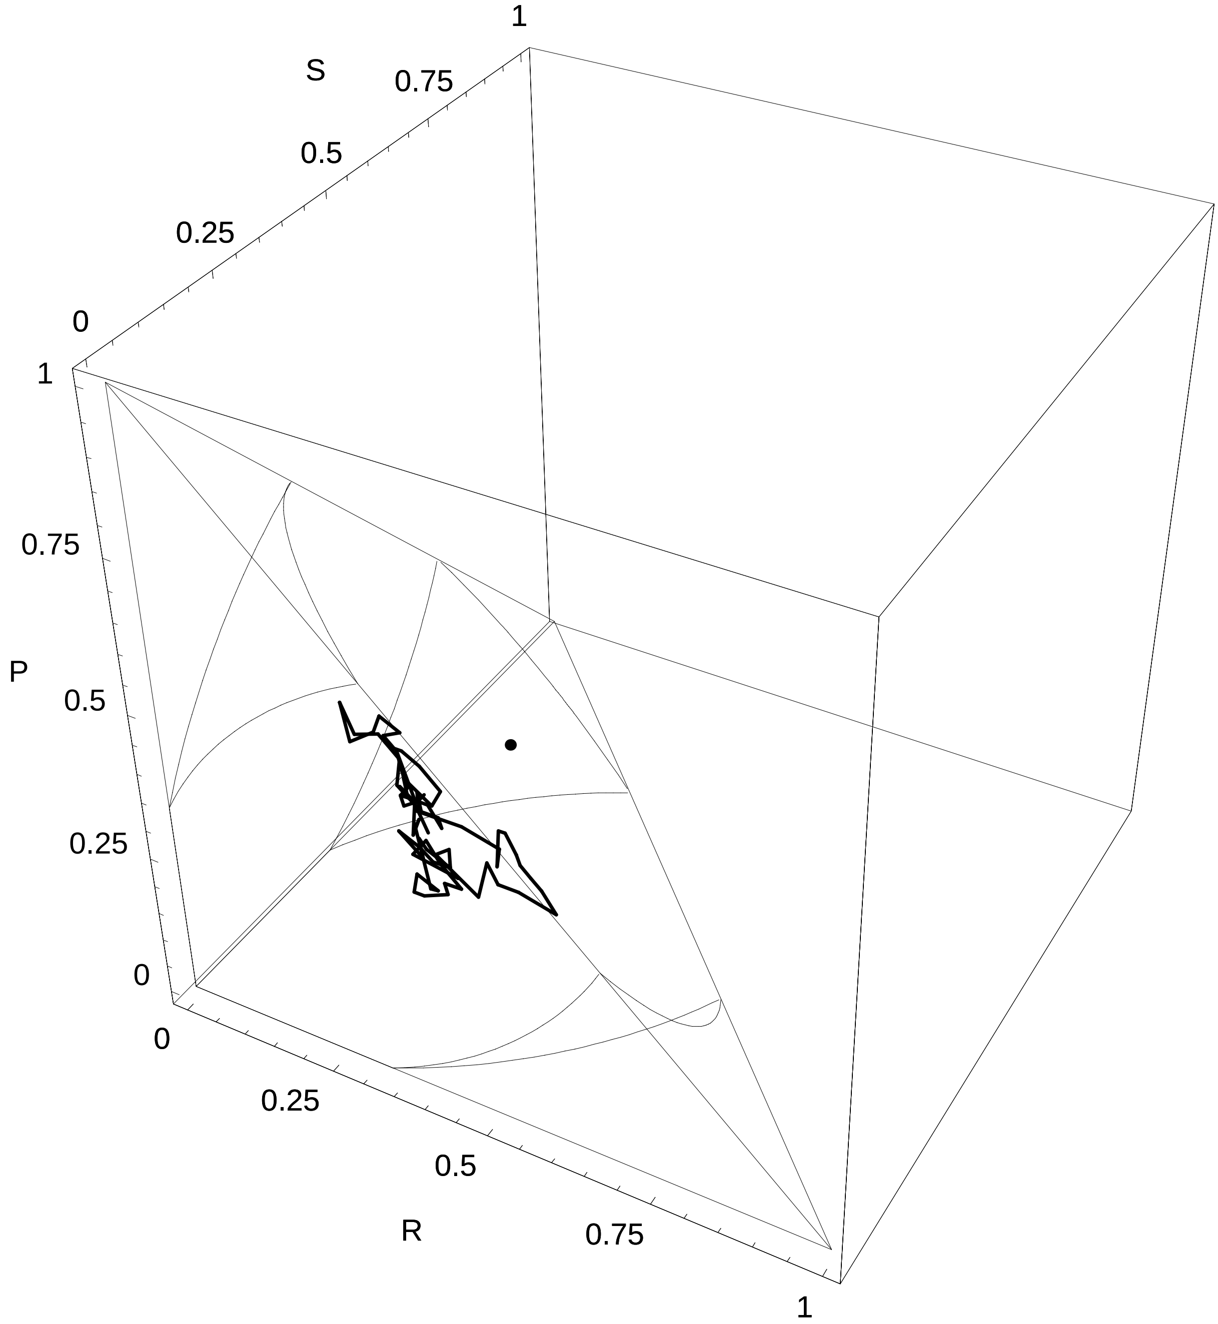

Figure 4: 10 Period Moving Average of Cohort 2 play path plotted in the simplex. Notice that bold line denoting data never penetrates wire frame polyhedra denoting regions of the state space in which *D* is an above average reply.

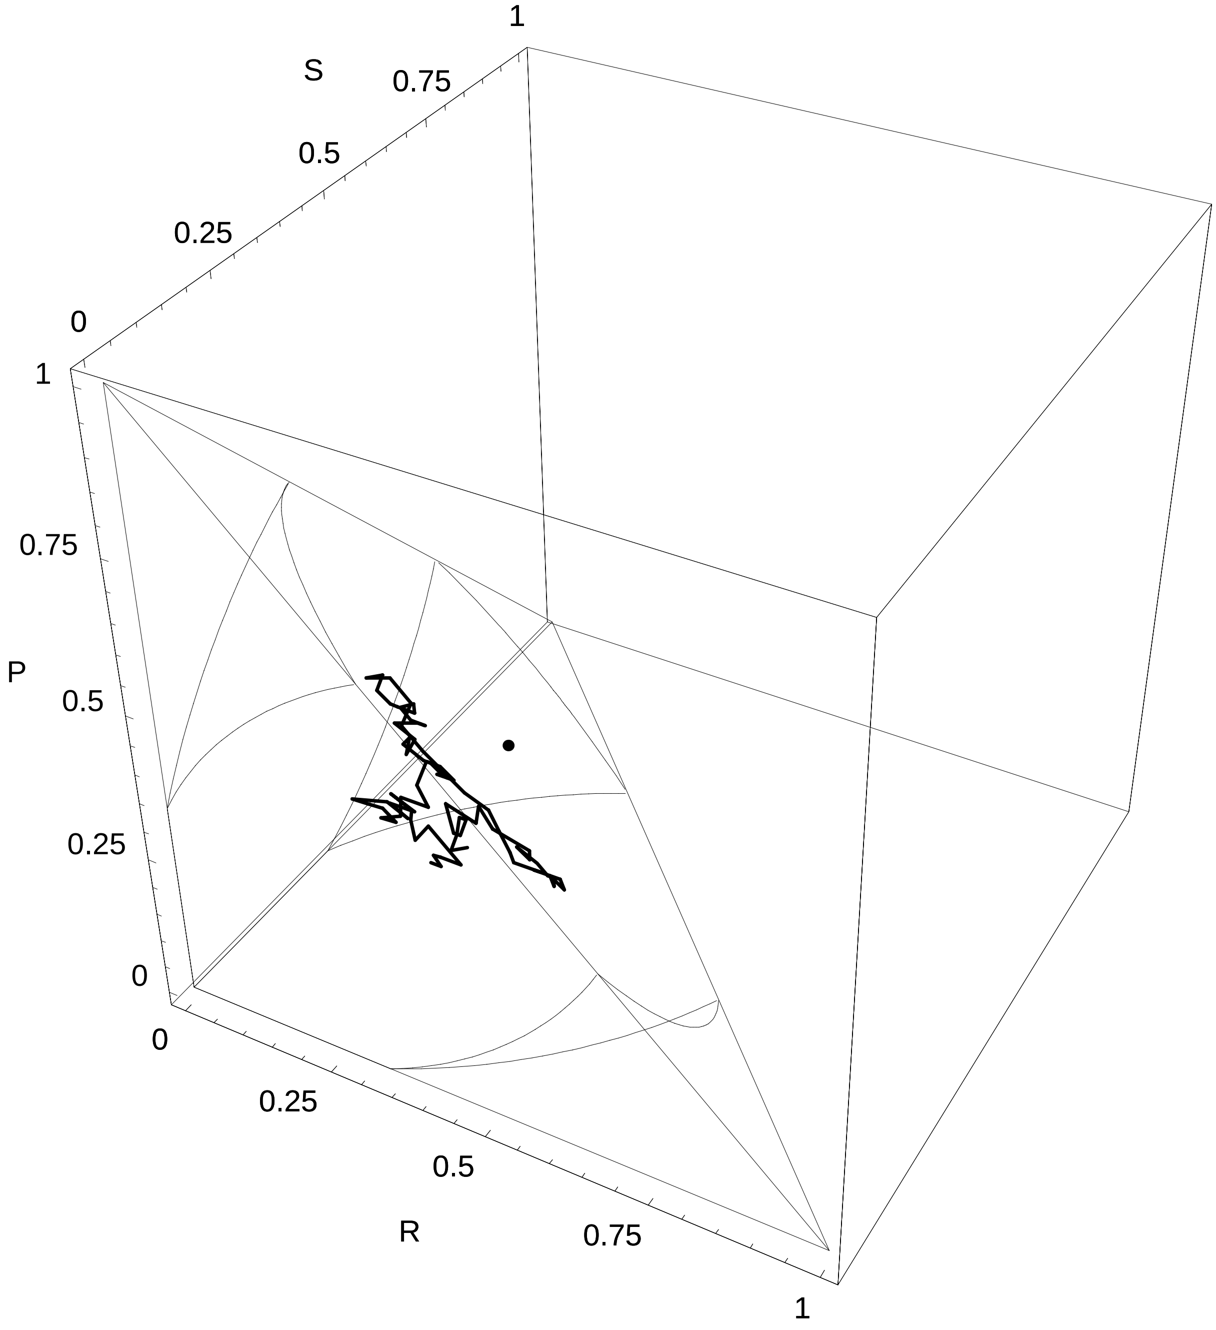

Figure 5 10 Period Moving Average of Cohort 3 play path plotted in the simplex. Notice that bold line denoting data never penetrates wire frame polyhedra denoting regions of the state space in which *D* is an above average reply.

'

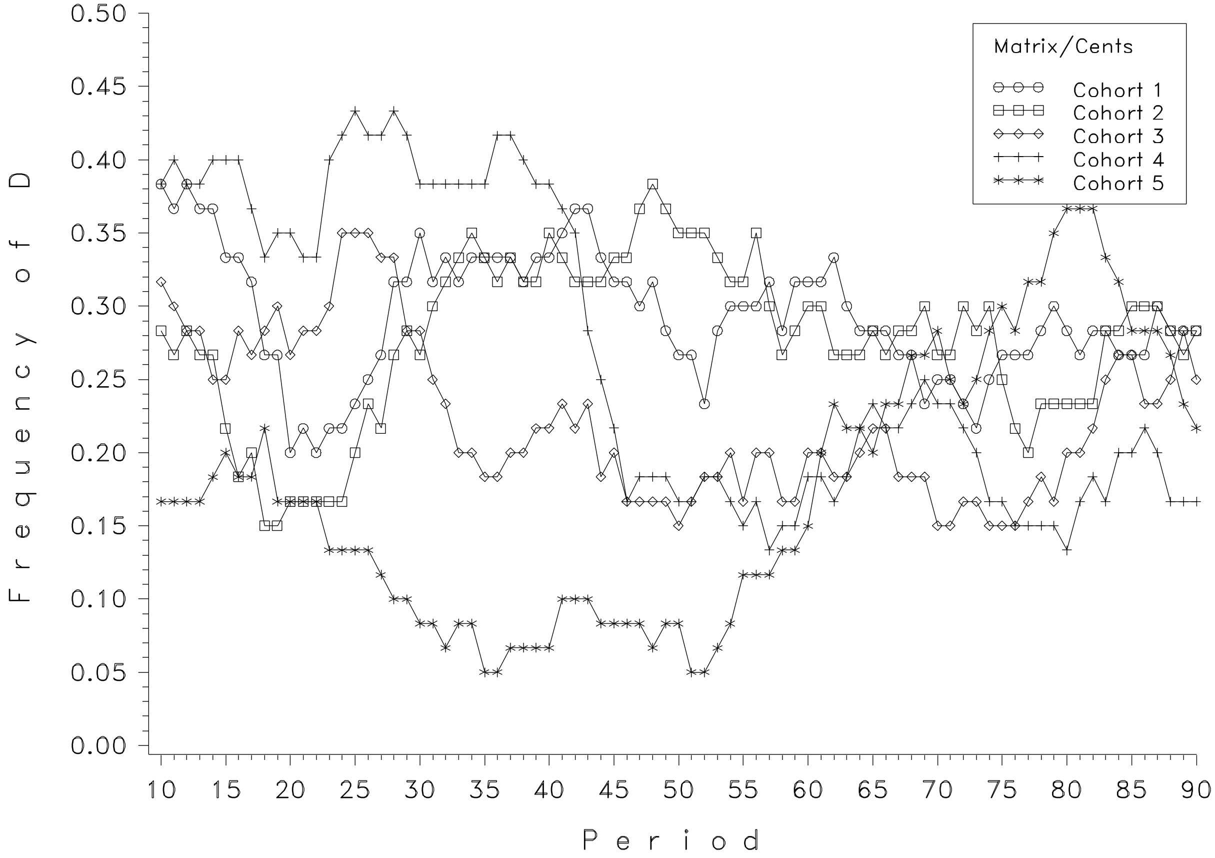

Figure 6: 10 Period Moving Average of *D* Frequency: Matrix/Cents Treatment.

'

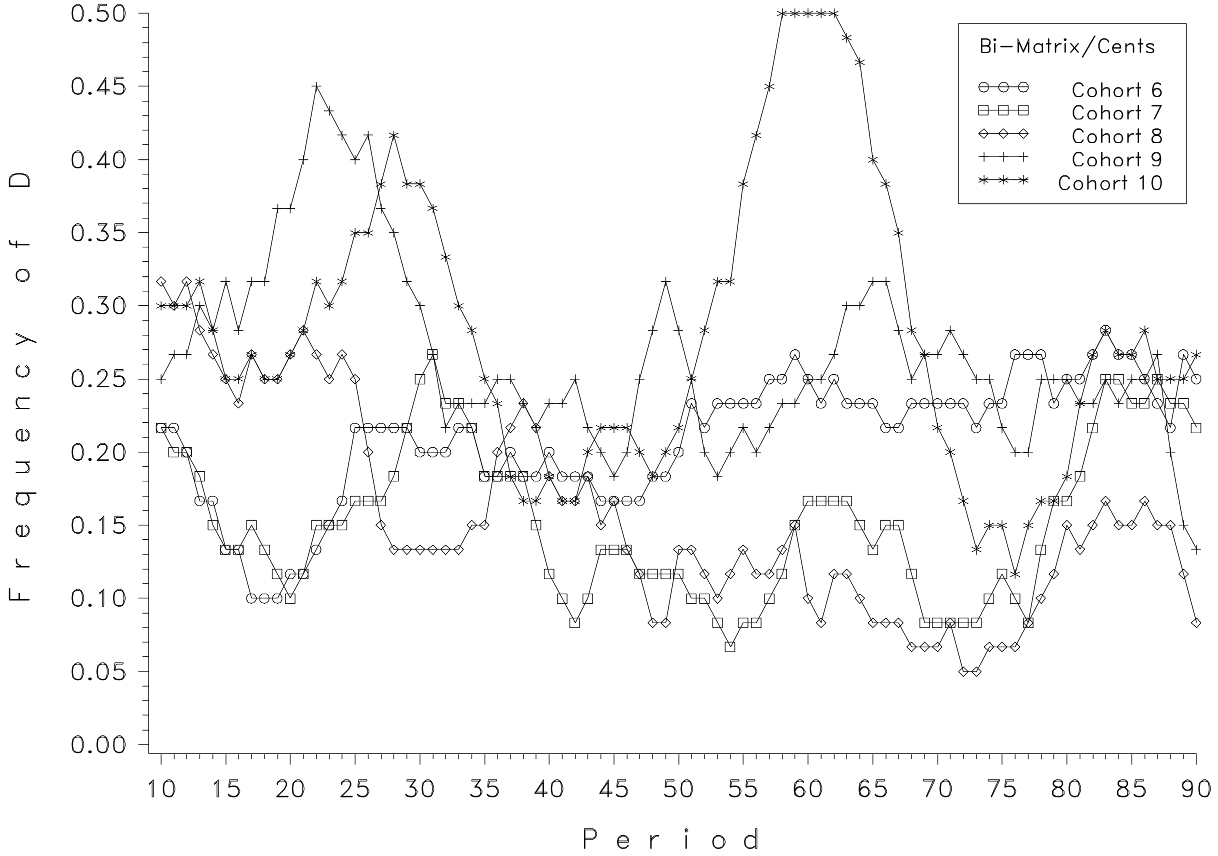

Figure 7: 10 Period Moving Average of *D* Frequency: Bi-Matrix/Cents Treatment.

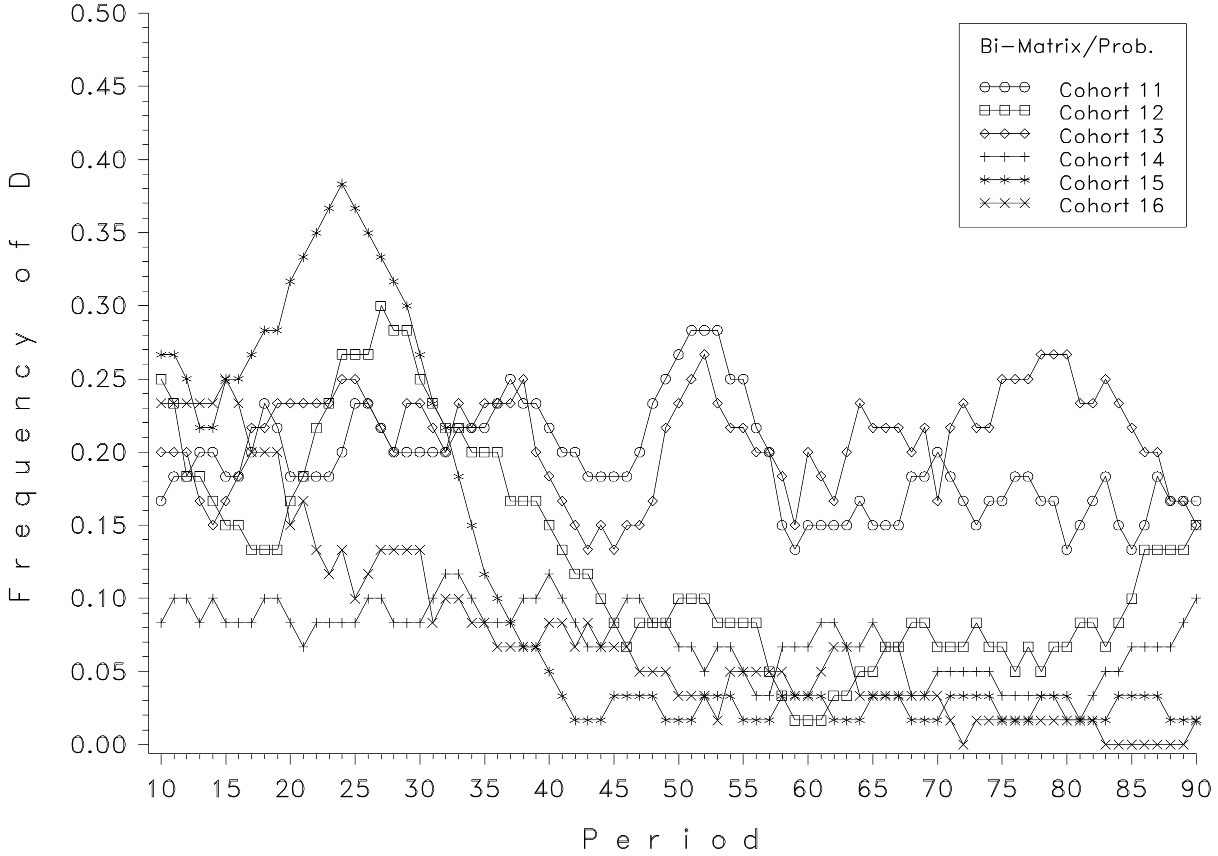

Figure 8: 10 Period Moving Average of *D* Frequency: Bi-Matrix/Probability Treatment.

'

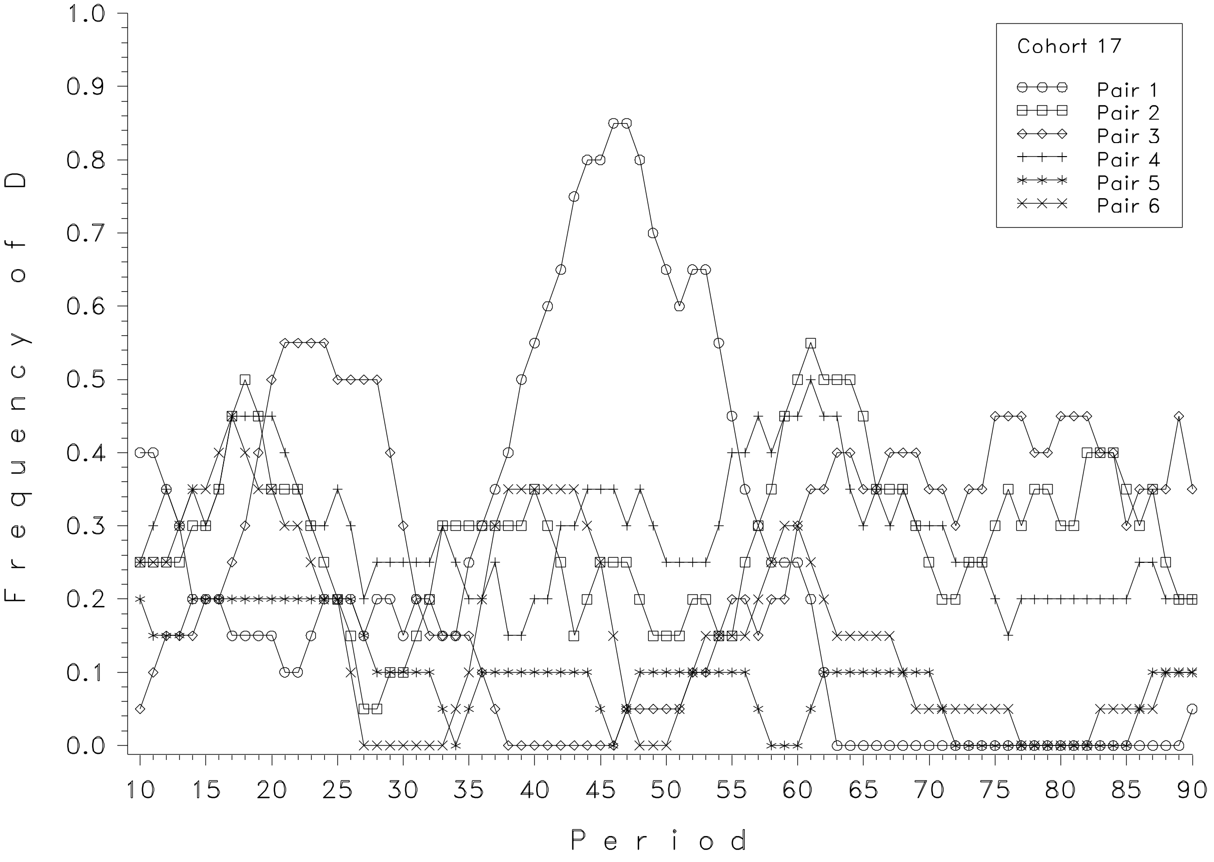

Figure 9: 10 Period Moving Average of *D* Frequency: Pairs Treatment: Cohort 17.

'

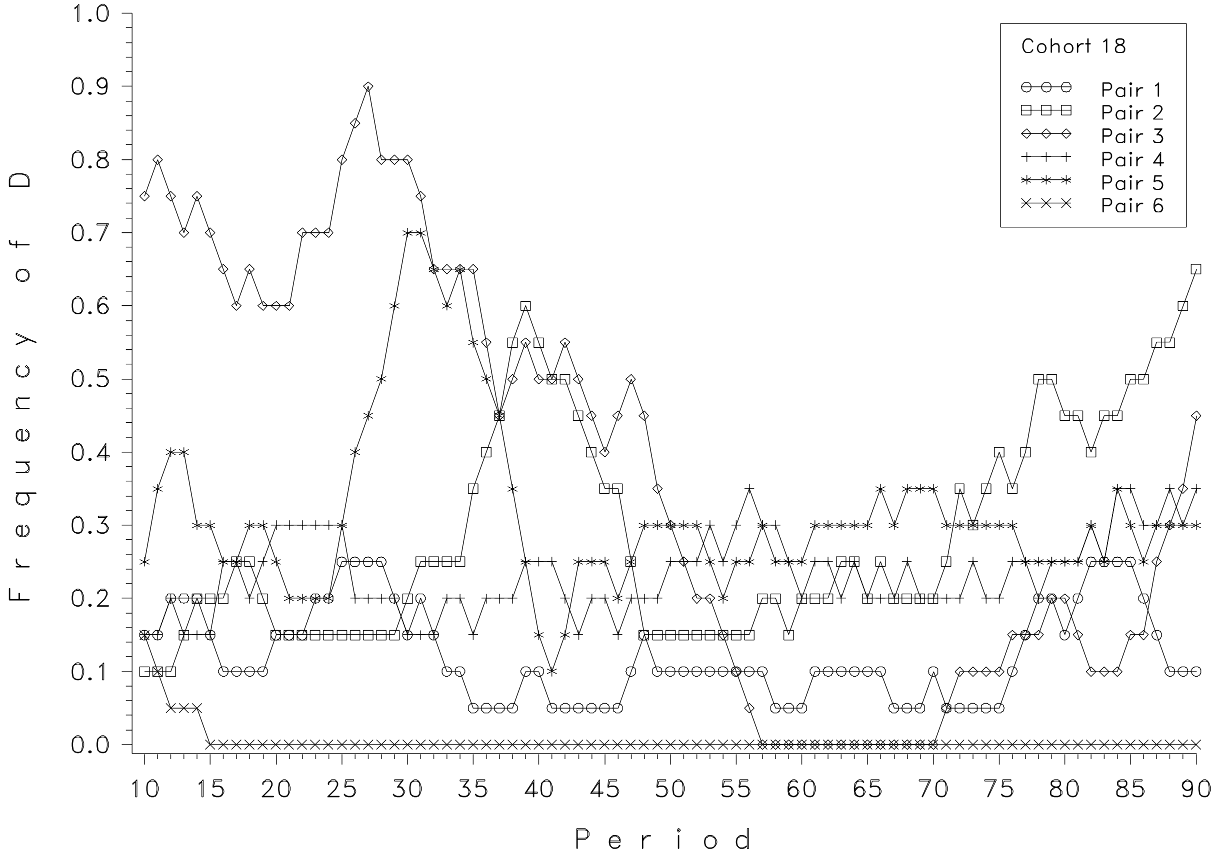

Figure 10: 10 Period Moving Average of *D* Frequency: Pairs Treatment: Cohort 18.