ARTICLE IN PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of Economic Theory ▮▮▮ (▮▮▮▮) ▮▮▮–▮▮▮

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Exploitable actions of believers in the "law of small numbers" in repeated constant-sum games

Steven Scroggin[1]

Department of Economics 0508, University of California San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0508, USA

Received 21 May 2004; final version received 7 November 2005

Abstract

In repeated fixed-pair constant-sum games with unique equilibria in mixed strategies, such as matching pennies, the subgame perfect equilibrium is repeating the stage-game mixed-strategy equilibrium action. In such games rational players avoid strategies that are *exploitable*, in that current actions either deviate systematically from the equilibrium action probabilities or fail to be serially independent of past actions. I revisit classic experiments and find that subjects' actions are sometimes exploitable because they are serially dependent. Subjects have difficulty in producing serially independent actions and in recognizing serially dependent sequences due to a bias called local representativeness.

JEL classification: C7; C9

Keywords: Behavioral economics; Experimental economics; Bounded rationality; Local representativeness; Game theory

1. Introduction

In repeated fixed-pair constant-sum games with unique equilibria in mixed strategies, the only subgame perfect equilibrium is repeating the stage-game mixed-strategy equilibrium action [19]. In such games rational players avoid actions that are *exploitable*, in that current actions either deviate systematically from the equilibrium action probabilities or fail to be serially independent of past actions.

In experiments with such games, subjects' actions are sometimes exploitable. Research on this issue has focused mainly on systematic deviations from the equilibrium action probabilities, but

E-mail address: sscroggi@weber.ucsd.edu.

[1] A version of this paper is a chapter in my dissertation.

0022-0531/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.jet.2005.11.001

there is reason to believe that serial dependence is also important. Cognitive psychologists who study individual decisions [4] and more recently some who study behavior in repeated fixed-pair games with mixed-strategy equilibria [13, "R&B"] question whether people can produce serially independent sequences when motivated to do so, and whether they can recognize serial dependence in others' actions.

Models of the effort to produce serially independent actions and recognize serially dependent sequences have their roots in Kahneman and Tversky's "representativeness" heuristic [7]. In particular, subjects often display a systematic bias known as "local representativeness", whereby they exaggerate how likely it is that a small sample resembles the parent population from which the sample is drawn. For example, the sequence HHHTTT is viewed by most people as more representative of repeated flips of an unbiased coin than the sequence HHHHHH, because the first sequence has a balanced number of H and T while the second sequence is all H. The sequence HTHTTH is viewed by most people as more representative of repeated flips of an unbiased coin than the sequence HHHTTT because the first sequence has more alternation (5 runs) than the second (2 runs).

Although data from experiments with repeated fixed-pair constant-sum games with mixed-strategy equilibria have been closely examined for exploitability, the learning models that have been used to detect it are not well suited to detect the kind of serial dependence that local representativeness implies. Learning models such as fictitious play impose linearity (and so additive separability) on players' responses to prior actions. By contrast, the usual models of local representativeness determine current actions as non-linear and non-separable functions of prior actions. Linear learning models may not reliably detect exploitability that stems from a source that appears plausible in light of the evidence from cognitive psychology.

In this paper, I revisit the issue of exploitability by reexamining the data from the classic experiments of Mookherjee and Sopher [8, "M&S"] and R&B. Because these experiments used matching pennies games with symmetric payoffs, they allow me to focus sharply on detecting non-linear serial dependence, while avoiding some of the difficulties associated with deviations from the equilibrium action probabilities. I model subjects' behavior via a class of simple dynamic decision rules that allow for the kinds of non-linear, non-separable responses to prior actions that local representativeness implies. These rules allow me to explore the extent to which subjects can produce sequences of actions with no readily perceived temporal pattern, and the extent to which subjects can recognize temporal patterns in their opponents' sequences.

Four decision rules defined precisely in Section 2 give rise to temporal sequences that might reflect representative behavior. The four decision rules I call representativeness ("Rep") [2] , best response to representativeness ("BR Rep"), counter representativeness ("C-Rep") and best response to counter representativeness ("BR C-Rep"). Two of these rules are *defensive* in that players following them try to be unexploitable by conditioning on their own actions. The other two, BR Rep and BR C-Rep, are *offensive* in that, players following them try to exploit their opponent's actions by conditioning on them. C-Rep is the best response to the best response to Rep, or action contrary to representativeness; the best responses to these four rules form a closed cycle (Table 1). [3]

Identification proceeds player-by-player and period-by-period by estimating the probability that a player is following a decision rule based on actions consistent with the rule in a set of recent

[2] This decision rule is related to a pattern explored by R&B.

[3] The decision rules have no known optimality properties. A theory of optimal or equilibrium actions by or against players subject to representativeness bias is the subject of work in progress with Vincent Crawford.

Table 1
A cycle of testable patterns

Pattern	Information set	Motivation	Relation to Rep.
Rep.	Own play	Defense	Rep.
BR Rep	Opponent play	Offense	BR(Rep.)
C-Rep.	Own play	Defense	BR(BR(Rep.))
BR C-Rep	Opponent play	Offense	BR(BR(BR(Rep.)))
Rep.	Own play	Defense	BR(BR(BR(BR(Rep.))))

overlapping sequences of actions. The identification strategy thereby allows for heterogeneity among players and for changes in a player's decision rules over time.

My data analysis uses one identification strategy for the defensive rules, Rep and C-Rep, and a related identification strategy for offensive rules, BR Rep and BR C-Rep. Each identification strategy is applied in two stages in each period. In the first stage of the identification strategy for a defensive rule at time *t*, a set of overlapping sequences is taken from the player's own prior actions and compared to the rule. If the set of sequences fits the rule with sufficient confidence, then the player's own prior actions are used to forecast the player's next action. If no forecast is warranted, I assume a hypothetical opponent trying to exploit the player obtains the 50% success rate that prior research suggests. By contrast, in the first stage of the identification strategy for an offensive rule at time *t*, each sequence in a set of overlapping sequences is assembled by concatenating a player's action in period $t - i, i > 0$ with a sequence of opponent actions prior to $t - i$. If the player wins by matching the opponent and the player is responding to what the player perceives as opponent actions according to a defensive decision rule, then the player will next act according the opponent's rule. If the set of concatenated sequences fits the rule with sufficient confidence, in the second stage a forecast about the player's period *t* action is made based on the rule.

After sufficient prior actions exist, the actions of any player and at any period may be identified as consistent with none, one or more of the rules. If actions are consistent with more than one rule, a forecast is made if a majority of the rules agree; otherwise, no forecast is made.

If a player's actions actually are serially independent, then sampling variation that by chance approximates a decision rule will pass the first stage. Forecasts based on actions which pass the first stage due to sampling variation will be insignificant. [4] If second stage forecasts are significant, then a hypothetical player acting on the forecasts would do better than the player could do with familiar alternatives, including equilibrium actions. In the M&S and R&B data sets, the forecasts in fact do perform substantially and significantly better than the familiar alternatives.

Typical results are given in Table 11, Column 1. In R&B 4018 (31.5%) overlapping sequences of actions of length 7 could be forecasted using the Rep decision rule with sufficient confidence to warrant an attempt to exploit them. The resulting forecasts were correct 57.8% of the time. The unconditional success rate, taking into account that only 31.5% of actions were found to be consistent with the Rep rule, was 52.4%. Two other rules, C-Rep and BR C-Rep, were also significant. BR Rep was sometimes found with sufficient confidence to attempt to exploit, but was not significant.

[4] Hereafter a decision rule is *significant* if the difference between forecasts of actions based on the rule and forecasts based on serially independent equilibrium action probabilities is statistically significant at the 95% level and the rule is more often correct.

Section 2 develops the econometric specification. Section 3 describes the experimental data of M&S and R&B. Section 4 describes the results of the present analysis. Section 5 is a discussion.

2. Testable implications of local representativeness

The literature contains several theories of how serially dependent sequences may arise from local representativeness. One theory of local representativeness uses balance. Another theory, alternation, is assessed by a runs test. [5] The following illustrates representativeness scores $br(\cdot)$, a weak lexicographic ordering over sequences formed by adding balance (times sequence length) to the number of runs.

2.1. Matching pennies example

Consider sequences of length l=3. Table 2 shows the 2^3 distinct sequences along with their balance, runs and representativeness score. The sequence $i_3 = (0, 0, 0)$ has score 1 and is relatively unrepresentative. On the other hand, sequence $i_3 = (1, 0, 1)$ has score 6 and is highly representative. The distinct sequences fall into three equivalence classes of scores. The representative ordering of scores for $l = 3$ is

$$w_6 Pr[br(i_3) = 6] > w_5 Pr[br(i_3) = 5] > w_1 Pr[br(i_3) = 1],$$

where w_k is a weight. [6] Here $w_6 = w_1 = \frac{1}{2}$, $w_5 = \frac{1}{4}$. The weights can be obtained by tabulating frequencies of scores reported in the last column of Table 2.

Now consider a thought experiment in which a player is serially independent for three rounds, but is representative in the fourth round of a series of non-overlapping sequences. Given three serially independent actions, the distribution of the eight distinct three-step sequences will be uniform, Table 2. If the player now chooses the representative action, the result is the set of eight four-step sequences given in Table 3. Table 4 sorts these eight sequences plus the eight other distinct four-step sequences into equivalence classes of scores and gives weights w_k.

The representative ordering of four-step scores is

$$w_{12} Pr[br(i_4) = 12] > w_{11} Pr[br(i_4) = 11] > w_{10} Pr[br(i_4) = 10] > w_7 Pr[br(i_4) = 7] \\ > w_6 Pr[br(i_4) = 6] > w_1 Pr[br(i_4) = 1].$$

---

[5] Balance, runs and representativeness scores are used informally in Section 1 and formally defined later in Section 2.2. A related approach, not pursued here, is the half-facetiously labelled “law of small numbers”. Believers in the law of small numbers act as if they believe that random sequences are actually produced by sampling *without* replacement from a finite set of possible outcomes. Rabin [12] explores the theoretical implications of the “small numbers” model for individual decisions.

[6] Each of the 2^l distinct sequences has one of not more than l^2 distinct scores. There is a many-to-one mapping from distinct sequences to distinct scores. Some scores may be more frequent because a larger number of distinct sequences are assigned that score. The weights w_k are a necessary adjustment that arises from differences in the number of distinct sequences assigned to a score. If k is the score of an equivalence class of sequences, let w_k be the inverse of the number of distinct sequences in the class. For representative actions, if x and y are two valid scores and $x > y$ then

$$w_x Pr[br(i_l) = x] > w_y Pr[br(i_l) = y],$$

where $Pr[\cdot]$ refers to frequencies and $Pr[br(i_l) = x]$ is the frequency of scores equal to x in sequences of length l.

Table 2
Three-step paths

Path i_3	Balance $b(i_3)$	Runs $r(i_3)$	Score $br(i_3)$
(0,0,0)	0	1	1
(0,0,1)	1	2	5
(0,1,0)	1	3	6
(0,1,1)	1	2	5
(1,0,0)	1	2	5
(1,0,1)	1	3	6
(1,1,0)	1	2	5
(1,1,1)	0	1	1

Table 3
Four-step paths in the thought experiment

Path i_4	Balance $b(i_4)$	Runs $r(i_4)$	Score $br(i_4)$
(0,0,0,1)	1	2	6
(0,0,1,1)	2	2	10
(0,1,0,1)	2	4	12
(0,1,1,0)	2	3	11
(1,0,0,1)	2	3	11
(1,0,1,0)	2	4	12
(1,1,0,0)	2	2	10
(1,1,1,0)	1	2	6

Table 4
Four-step representative equivalence classes

Score $br(i_4) = k$	Weight w_k
1	1/2
6	1/4
7	1/4
10	1/2
11	1/2
12	1/2

Rank correlation of the three-step sequences is zero by construction, but when the next action is representative, the rank correlation of the sequence frequencies becomes notable. By inspection, 10 of the 15 inequalities in the representative ordering hold for the data in Table 3; there are 4 ties and 1 inequality fails. As an illustration of one of these inequalities consider scores 7 and 12. In Table 3 no sequences have score 7, so the weight-adjusted frequency of scores of 7 is zero. Two sequences have score 12, with weight $1/2$, so the weight adjusted frequency of scores of 12 is 1.

Table 5
Rank correlations in the thought experiment

Rank correlation $r(\cdot)$	Value
ρ	0.61
$\hat{\rho}$	0.66
τ	1.20[a]
$\hat{\tau}$	0.87

[a]The tie-adjustment formulas do not confine reported statistics to $[-1, 1]$, as illustrated here for Kendall's τ with tie adjustment. Hereafter, except as noted, I use p-values associated with correlations rather than the correlations themselves.

Since

$$w_{12}Pr[br(i_4) = 12] = \left(\frac{1}{2}\right)(2) > w_7Pr[br(i_4) = 7] = \left(\frac{1}{4}\right)(0),$$

the representative ordering holds for scores 7 and 12 in this hypothetical data.

Table 5 reports rank correlations for the thought experiment. Since the sequences do not overlap, I cannot assess p-values for these rank correlations using the technique I use in the remainder of the paper; however, all the rank correlation measures are highly positive.

2.2. Econometric model

I now translate the balance and runs notions that underlie local representativeness into behavioral assumptions that apply to repeated fixed-pair matching pennies games. Although it is not hard to envision extensions to games with more than two actions, and more than two players, I state the assumptions here only for the symmetric matching pennies games used to generate the data I analyze. Scroggin [16] discusses more general settings.

Let there be 2 matching pennies players and let $s_{i,t}$ be player i's action ($i = 1, 2$) in round t from the binary action space $S = \{1, 0\}$. Player i's l-step *sequence* i_l is an element of the l-tuple of sets $L = S^l$, specifically, the time-ordered list of i's most recent actions,

$$i_{l,t} = (s_{i,t-l+1}, s_{i,t-l+2} \ldots s_{i,t}).$$

Hereafter, I suppress the t subscripts when referring to sequences and their elements to simplify the notation, and

$$i_l = (s_{i,1}, s_{i,2} \ldots, s_{i,l}).$$

Players view as more balanced, hence more "random", a sequence which is closer to that of the population in cumulative (within the sequence) frequencies of actions. Define the *balance* function $b : L \to N$ as

$$b(i_l) = l - \max_{s \in S} \sum_{k=1}^{l} I(s_{k-l} = s),$$

where $I(\cdot)$ is the indicator function and N is the set of natural numbers (including zero).

Sequences with more runs are perceived as more “random”. Define the *runs* function $r : L \to N$ as

$$r(i_l) = 1 + \sum_{k=1}^{l-1} I(s_{i,k-l} \neq s_{i,k-l+1}).$$

This follows the standard definition. [7]

Both balance and runs are ordinal; representativeness is a lexicographic weak ordering of sequences by balance, breaking ties in balance with runs. Define the *representativeness function* $br : L \to N$ as

$$br(i_l) = lb(i_l) + r(i_l)$$

and call $br(i_l)$ the representativeness score (or *score*) of sequence i_l. Relatively representative sequences have relatively high scores.

A player produces $t - l + 1$ overlapping sequences of length l up to time t. Each sequence produced is one of 2^l distinct sequences. For serially independent actions, each distinct sequence is equally likely. A 2^l-way tie is the ordering of sequence frequencies under the null hypothesis. The alternative hypothesis is the representative ordering of sequence frequencies. A set of sequences is representative if the rank correlation between its sequence frequencies and the representative ordering of sequence frequencies is significantly greater than zero.

There are two standard rank correlation estimators, Kendall's $\hat{\tau}$ and Spearman's $\hat{\rho}$. I use both, as well as versions of each that adjust for ties in rank, τ and ρ. A researcher's choice among these four measures makes little difference. Formulas for tie adjustments are given in [16, p. 61]. Critical values for overlapping sequences are obtained using a parametric bootstrap technique. That also is discussed in [16, p. 61].

Forecasting is next. Suppose, in choosing her next action, player i conditions on the prior actions of exactly one player j. If $j = i$ then player i conditions on her own prior actions, and plays defense; if $j \neq i$ then player i conditions on the prior actions of another player and plays offense. Let j_l be the sequence consisting of the last l actions. Define $j_l|s_i$ to be the concatenation of j_l and the action s_i, which is a sequence of length $l + 1$. Then

$$\hat{br}(j_l|s_i) = \underset{s_i \in S}{\operatorname{argmax}}\ br(j_l|s_i)$$

gives the current action s_i which maximizes the representativeness score conditional on the last l actions. If $j = i$, $\hat{br}(j_l|s_i)$ is the *representative forecast*. If $j \neq i$, and i wins by matching with j, $\hat{br}(j_l|s_i)$ is the best response to representativeness in j.

The econometric models take the form

$$s_{i,t} = f_L(I(r(j) > x)\hat{br}(j_l|s_i)\beta + X_i'\gamma + \varepsilon_{i,t}), \quad i = 1, \ldots, I, \quad t = \underline{w} + l, \ldots, T.$$

The dependent variable $s_{i,t}$ is player i's action at time t. Since $s_{i,t}$ is binary-valued, I use a logit regression $f_L(\cdot)$. Here i is an index to the I individual subjects in a data set. The coefficient of interest β estimates the coefficient for a conditional forecast.

[7] The notion that the number of runs measures perceived randomness is not persuasive at the extreme of maximal runs: simple alternation is not perceived to be random. A useful extension of this paper would be a better measure of perceived randomness associated with alternation.

The condition is expressed by the indicator function $I(r(j) > x)$. The indicator is 1 if a subset of prior overlapping sequences of length l of player j using rank correlation estimator $r(\cdot)$ has a value which exceeds the critical value associated with the threshold probability x. The subset of prior overlapping sequences assessed by the indicator is the lesser of the $\overline{w}$ most recent sequences and all the prior sequences that exist, but at least $\underline{w}$ sequences. If the indicator is 0 or if less than $\underline{w}$ sequences exist, the indicator fails and the observation is removed from the regression. If the indicator is 1, then a representativeness forecast $\hat{br}(j_l|s_i)$ is active. Player fixed effects are given by X_i; γ is a vector of fixed effects coefficients;[8] and $\varepsilon_{i,t}$ is an error term.

Aside from robustness checks, $\underline{w} = 1$, $\overline{w} = 50$, $x = 0.8$ and $r(\cdot) = \rho(\cdot)$.

If $j = i$ we test Rep and the coefficient is labelled β_1, if $j \neq i$ we test whether i has a history of best responding to representativeness BR-Rep in j and the coefficient is labelled β_2.

C-Rep is negative rank correlation. Because counter-representative sequences lack balance and runs, they tend to be streaky relative to serially independent sequences. The C-Rep specification is

$$s_{i,t} = f_L(I(r(j) < -x)(1 - \hat{br}(j_l|s_i))\beta + X_i'\gamma + \varepsilon_{i,t}),$$
$$i = 1, \ldots, I, \quad t = \underline{w} + l, \ldots, T, \tag{1}$$

If $j = i$ we test C-Rep with β_3; if $j \neq i$ we test best responses to counter-representativeness in j BR C-Rep with β_4.

One can linearly combine any permutation of these models and get a new model. The most general is the portmanteau model:

$$\begin{aligned} s_{i,t} = f_L \Big[& I(r(i) > x)\hat{br}(i_l|s_i)\beta_1 \\ & + I(r(j) > x)\hat{br}(j_l|s_i)\beta_2 \\ & + I(r(i) < -x)(1 - \hat{br}(i_l|s_i))\beta_3 \\ & + I(r(j) < -x)(1 - \hat{br}(j_l|s_i))\beta_4 \\ & + X_i'\gamma + \epsilon_{i,t} \Big], \quad i = 1, \ldots, I, \quad t = \underline{w} + l, \ldots, T. \end{aligned}$$

Here I assume $j \neq i$ and so β_1 is the coefficient on Rep forecasts, β_2 is BR-Rep, β_3 is C-Rep and β_4 is BR C-Rep.

Forecasts from components of the portmanteau model might conflict; however, most conflicts are avoided by construction. Except for robustness checks, the correlation threshold x is always set to a positive value. If the correlation threshold x is positive, then at most one of the defensive models (Rep and C-Rep) can be found to be likely enough to merit a second-stage forecast at a particular t; similarly, at most one of the offensive models (BR Rep and BR C-Rep) will merit forecasting. It remains possible for offensive and defensive strategies to produce contradictory forecasts; however, contradictions are not a problem in practice (see Table 11, Panel B, Column 1 and related text).

3. Data

M&S and R&B report experiments and the data are from those experiments.[9] Both experiments were repeated fixed-pair symmetric constant-sum games like the school yard version of matching pennies, Fig. 1.

---

[8] In most of the regressions a constant is used because fixed effects were not significant.

[9] I thank both Barry Sopher and David Budescu for their data.

	0	1
0	0 1	1 0
1	1 0	0 1

Fig. 1. Matching pennies normal form game.

The M&S and R&B experiments are well suited to the study of representativeness. Matching pennies is a game of pure competition. There is no element of coordination and so no complication from attempts to coordinate, as in [17]. In both data sets, players typically take pure actions with frequencies that closely approximate over time the equilibrium action probability. Matching pennies is the simplest game in which representativeness plays an important role. [10]

The M&S design produces positive expected payoffs for both players. Subjects were economics students at the University of Delhi. Ten pairs of players played 40 rounds of two treatments. I use only Treatment 2. [11] R&B (Treatment D, "Dyad") was a psychology experiment. Forty-five pairs of undergraduates from the University of Haifa played 150 rounds of matching pennies. [12]

In both experiments, a repeated fixed-pair time series of row player and column player actions comes from Row's perspective. Column is isomorphic to Row, and by transforming the data—by reversing Column's actions (assigning 0 to 1 and 1 to 0) and assigning the Row label to Column and vice versa—a second repeated fixed-pair time series is available. Each subject appears in exactly one series as Row and in exactly one other series as Column. The connection between the two series is ignored hereafter.

Both papers contain descriptive statistics which I do not repeat, but I do report one new tabulation. I consider the set of all non-overlapping sequences in the R&B data. For sequences of selected even lengths, Table 6 reports the number of such sequences, the number that are balanced (equal number of heads and tails), the number of streaks (all heads or all tails) and the rest. Along with these sequence frequencies, I report the ratio between the sequence frequencies and expected number of sequences for each type of sequence. The expectation is based on the null hypothesis that actions are serially independent, 50% heads. For example, there are 1504 sequences of length 8, 542 of which are balanced and 7 are streaks, 1.31 and 0.6 times (respectively) as frequent as one would expect in the limit in serially independent actions.

A referee suggested a simple alternative hypothesis to the full representativeness model. Perhaps subjects switch actions more than half of the time. Indeed, R&B subjects do switch actions 52.6% of the time. However, to produce the proportion of balanced sequences seen in the data for length l=4, the switching probability would be 57.7% and to produce the proportion of balanced

[10] There may be many games in which players commonly fail serial independence and fail to match their actions to the equilibrium action probability. This is an area for future work.

[11] Treatment 2 was matching pennies; in Treatment 1 subjects were not told the payoff matrix, opponent's actions or opponent's realized payoffs.

[12] Subjects received an initial endowment of 20 New Israeli Shekels. Some lost their endowment and the game terminated. Consequently, the data is an unbalanced panel. R&B had two other treatments that had no strategic interaction.

Table 6
R&B data: balanced, streaky and other play

Path length	N	Balanced	Balanced/ E[balanced]	Streaks	Streaks/ E[streaks]	Rest	Rest/ E[rest]
2	6184	3251	1.05	2933	0.94	—	—
4	3088	1348	1.16	316	0.81	1414	0.92
6	2010	794	1.26	40	0.64	1176	0.89
8	1504	542	1.31	7	0.60	955	0.88
12	1000	300	1.32	1	2	699	0.90
20	582	159	1.55	0	—	423	0.88

sequences seen in the data for length $l = 6$, the switching probability would be 61.2%. The data are not consistent with simple switching.

4. Results

The first test is for representativeness in M&S aggregated across subject pairs and time.[13] Table 7 gives rank correlations in sequence frequencies from the M&S data. Kendall's τ and Spearman's ρ (both with tie adjustment) are reported, with sequence lengths from 4 to 9.[14]

All of the reported statistics are significant at least at the 10% level. If sequences are representative at length $l = z$ they may be representative at $l = z + 1$ since the sequences differ only in the $z + 1$th element. So the tests reported in Table 7 are not independent, but neither are they redundant; they are mutually supportive in that they all have the right sign and similar levels of statistical significance. This is evidence for the representative actions in the M&S data as a whole. I now turn to player-by-player turn-by-turn forecasting models.

Table 8, Column 1 gives results for the baseline Rep specification in M&S. Sequence length is $l = 5$; a single constant is estimated instead of pair fixed effects. With the first stage threshold critical value x set at 80% probability, 115 out of 680 observations, involving eight of the 20 players, passed the first stage. In the absence of a consensus measure of goodness-of-fit for logit models [6, p. 831], I use three measures: (1) percent correct, conditional on a forecast (sometimes called "efficiency"); (2) net wins; and (3) $\hat{\beta}$ z-score. Percent correct has a transparent intuition: 50% is worthless, 100% is perfect. The representativeness forecast was right 66.1% of time, 37 times more than it was wrong. The z-score on $\hat{\beta}_1$ is 3.38, p-value is 0.0004.[15] Rep is both a statistically and substantively significant forecaster of actions in M&S.

Table 8, Column 2 replaces the constant with pair fixed effects. An F test for the fixed effects as a whole was not significant. A constant is used instead of fixed effects hereafter.

---

[13] I analyzed M&S completely before turning to R&B to minimize data mining issues; however, the results for R&B are more significant and show more strategic behavior.

[14] Sequences shorter than $l = 4$ are difficult to analyze because there are only eight distinct sequences of length 3 (see Table 2). For sequences of length $l = 9$, there is plenty of theoretical variation—$2^9 = 512$ distinct sequences—but the frequencies may be poorly estimated for lack of data. In M&S expected sequence frequency for sequence length $l = 9$ is about 1.2.

[15] In this regression $\hat{\beta}_1 = 1.33$. This coefficient is the log of the odds ratio of the forecast. These coefficients are not reported hereafter since the sign and statistical significance of the coefficient along with efficiency and net wins measures are more easily interpreted.

Table 7
Representativeness in M&S data

Path lengths l	Kendall's τ	Spearman's ρ
4	0.64*	0.73*
5	0.47*	0.60*
6	0.38*	0.49*
7		0.39*
8		0.32**
9		0.24**

*10%, **5%.

Table 8
M&S data, $l = 5$

	Rep.	Rep.	Rep.	Portmanteau	Rep, Lags
Path length l	5	5	5	5	5
Fixed effects	No	Yes	No	No	No
Function	Logit	Logit	OLS	Logit	Logit
Forecasts	115	115	115	271	
Pairs used	8	8	8	18	
Percent correct	66.1	66.1	66.1	59.8	
Net wins	37	37	37	53	
$\hat{\beta}_1$ z-score	3.38	3.57		3.42	3.25
$\hat{\beta}_1$ t-stat			3.77		
Probability $\beta_1 = 0$	0.0004	0.0002	0.00008	0.0003	0.0005
$\hat{\beta}_2$ z-score				0.46	
$\hat{\beta}_3$ z-score				0.94	
$\hat{\beta}_4$ z-score				1.64	
$C_{i,t-1}$ z-score					1.66
$C_{i,t-2}$ z-score					−1.52
$C_{i,t-3}$ z-score					0.05
$C_{i,t-4}$ z-score					0.11

Now I vary parameter values to explore robustness. First, I vary the sequence length, adjusting the first stage threshold critical value x to maintain 80% probability of the Rep decision rule. This is summarized in Table 9. The results are significant for all sequence lengths. The first and second best on all three of the ranking measures—percent correct, net wins and z-score—are odd-length sequences. Odd-length sequences cannot have ties in balance so odd-length sequences doing better is consistent with the runs test being a poor tie-breaker.

Next, I vary the first stage threshold critical value x. Only 51% of forecasts are correct with no first stage, but the percent correct increases smoothly to 77% for $x = 0.80$. The z-score varies from 0.54 (no first stage) to 3.87 ($x = 0.50$) and then falls to 2.76 ($x = 0.80$). One can obtain non-significant results for $x < 0.20$. This is evidence that some actions do not follow the Rep decision rule and a first stage is necessary to screen out these cases.

Next, I vary the rank correlation statistic: there are four to choose from and they are all significant. No one measure dominates.

Table 9
M&S path length l variations

	Rep.	Rep.	Rep.	Rep.	Rep.	Rep.	Rep.
Path length l	3	4	5	6	7	8	9
Forecasts	19	129	115	143	143	137	116
Pairs used	1	9	8	8	8	9	8
Percent correct	78.9	61.2	66.1	60.1	63.6	59.1	60.3
Net wins	11	29	37	29	39	25	24
$\hat{\beta}_1$ z-score	2.33	2.52	3.38	2.41	3.22	2.12	2.22
Probability	0.01	0.006	0.0003	0.008	0.00006	0.017	0.013

I ran more robustness tests by changing the conditioning window, parameterized by $\underline{w}$ and $\overline{w}$. In the baseline specification, Table 8, Column 1, changing $\underline{w}$=1 to $\underline{w}$=15 reduces the z-score for the representativeness variable from 3.38 to 3.20. Decreasing $\overline{w} = 50$ to $\overline{w} = 9$ in the baseline specification reduced the z-score from 3.38 to 3.34. The algorithm requires about $l+9$ consecutive actions to make optimal (as a function of the number of periods analyzed) forecasts. If players change decision rules too frequently, there will be insufficient data to detect their rules before they change them.

Here is a list of the parameters that one could manipulate in search of a result like the baseline representativeness specification, Table 8, Column 1: sequence length l, first stage threshold critical value x, correlation statistic $r(\cdot)$, minimum window $\underline{w}$, maximum window $\overline{w}$, fixed effects (yes/no), logit vs. OLS, and permutations of linear combinations the Rep decision rule with other rules (reported later). Without exception, the Rep results are robust to modest changes. One might also manipulate the data directly, omitting troublesome players, games or time periods. Here all the data were used. Or, one might use a different sequence length l for the first stage test than for the second stage forecast. I used the same l in both stages.

As another check, I simulated a serially independent data set using the bootstrap procedure described in [16] (but only one replication). I tried 10 regressions on this pseudo-random data set with different parameters for the Rep model. The regression with the best $\hat{\beta}$ had p-value 11%, though the sign was wrong.

Having found Rep in M&S, I turned to the three other decision rules. The portmanteau model is typical, Table 8, Column 4. BR Rep β_2, and C-Rep β_3, are not statistically significant. BR C-Rep, β_4, is barely statistically significant, but the result is not robust.

Table 8, Column 5 shows Rep combined with lags of opponent's actions. [16] Opponent's lags one and two are not significant alone (M&S have this result), but opponent's lags one and two are statistically significant, with opposite signs, when combined with the Rep model. The result is robust.

The details of the specification developed in M&S were applied to the R&B data without modification. The first stage threshold critical values x are specific to the number of games and their length in M&S. Nevertheless, for consistency, I retained the same critical values.

Results for Rep in R&B data are given in Table 10, Column 1. There were 3558 forecasts with z-score 9.46, probability 0.0000000. Column 2, BR Rep, is not significant and has the wrong sign.

[16] No results are reported for success measures because I have no theory for how to assess success for Rep with lags.

Table 10
R&B data, path length $l = 5$

	Rep.	BR(Rep.)	C-Rep.	BR(C-Rep.)
Path length l	5	5	5	5
Forecasts	3558	3630	1158	1129
Net wins	568	−84	54	85
Percent correct	58.0	48.8	52.3	53.8
$\hat{\beta}$ z-score	9.46	−1.44	1.56	2.46
Probability	0.0000000	0.07	0.058	0.007

Table 11
R&B data, path length $l = 7$

Panel A	Rep.	BR(Rep)	C-Rep.	BR(C-Rep)
Path length l	7	7	7	7
Forecasts	4018	4070	1194	1138
Percent correct	57.9	49.3	55.9	53.0
Net wins	632	−60	142	68
$\hat{\beta}_1$ z-score	9.92	−0.95	3.98	2.01
Pr ($\beta_1 = 0$)	0.0000000	0.17	0.00003	0.02

Panel B	Rep.,C-Rep. BR(C-Rep.)	plus lag
Path length l	7	7
Forecasts	5576	
Percent correct	57.2	
Net wins	806	
$\hat{\beta}_1$ z-score	9.97	9.96
Pr ($\beta_1 = 0$)	0.0000000	
$\hat{\beta}_3$ z-score	4.00	3.97
$\hat{\beta}_4$ z-score	2.21	1.83
$\hat{C}_{i,t-1}$ z-score		−3.11

Column 3, C-Rep, just misses statistical significance. Column 4, BR C-Rep is significant, but not nearly as salient as Rep.

The R&B game is almost four times longer than the M&S game, for those who went the distance. Longer sequences imply more distinct sequences, and allowing the sequence space to ramify is more promising in a larger data set. Table 11, Panel A is the same as Table 10, but with sequence length $l = 7$. Rep is as prominent as before and BR Rep retains the wrong sign and remains insignificant; however, the borderline results for C-Rep improve to strong statistical significance. BR C-Rep remains significant. Table 12 repeats the exercise for sequence length $l = 9$.

The three significant $l = 7$ decision rules, Rep, C-Rep and BR C-Rep are combined in Table 11, Panel B, Column 1. When forecast separately, 5212 forecasts are made in defensive

Table 12
R&B data, path length $l = 9$

	Rep.	C-Rep.	BR(C-Rep.)
Path length l	9	9	9
Forecasts	4210	1147	1067
Percent correct	57.4	59.8	53.1
Net wins	622	225	67
$\hat{\beta}$ z-score	9.54	6.48	2.02
Probability	0.0000000	0.0000000	0.02

strategies (Panel A, Columns 1 and 3) and 1138 in offensive strategies (Panel A, Column 4). When offensive and defensive strategies are combined, the number of forecasts is less than the sum of the parts, because there are cases in which there is both an offensive and a defensive forecast. Testing offense and defense together increases the number of forecasts over either alone, and does not affect forecast quality.

Table 11, Panel B, Column 2 adds the first lag of opponent's actions to Column 1. The first lag is significant. Longer lags were insignificant (not reported).

5. Discussion

At times players choose exploitable actions. The strongest results were for sequence length $l = 7$ in the R&B data, Table 11 and length $l = 5$ in the M&S data, Table 8, Column 1.

A player attempting unexploitable defense may be representative instead. Rep was prominent in both of the data sets. A player may play offense, recognizing her opponent's Rep decision rule and best responding. However, BR Rep was not found, perhaps because a typical player finds the Rep decision rule difficult to distinguish from serially independent behavior.

C-Rep may represent faulty defense. Alternatively, a player may anticipate BR Rep, whether or not it is present, with C-Rep. Whatever the source, C-Rep was significant in R&B data.

BR C-Rep was significant in R&B and marginal in M&S. BR C-Rep was sometimes *more* prominent than the C-Rep to which it is the best response. If a player adopts the C-Rep decision rule or produces a serially independent sequence that looks counter-representative due to sampling variation, her opponent may be extremely sensitive to it, since there may be more difference between the opponent's model of unexploitable sequences and C-Rep than there is between serially independent sequences and C-Rep. Further, having perceived counter-representative defense, the best response is deterministic. When the player responds with high probability to a subtle, or non-existent, signal, it may be easier to detect the signal plus response than the signal plus the next signal because the signal plus response is more consistent with the decision rule.

Any of these decision rules may be thought of as sophisticated—all can be interpreted as a response to an opponent.

The econometric models contain parameters that determine how to identify a decision rule, but the rules themselves are non-parametric. Since the decision rules are non-parametric, estimation and calibration errors described by Wilcox[17] are absent.

[17] N.T. Wilcox, "Heterogeneity and Learning Principles", Unpublished (2003).

The framework of decision rules plus two-stage identification is flexible but harbors structure. A reduced form identification strategy that used all the information in the information set would use all the prior actions of both players as conditioning information, with dummy variables for each distinct sequence. Even for matching pennies that requires 4^n dummy variables, where n is the number of rounds. I cut this down in several steps: (1) I assume memory is limited so only the past l actions are recalled, where $l < 10$. (2) I assume decision rules that allow conditioning on only one player's prior actions in one consistent way, though players may use more than one decision rule. (3) I suppose that sequences are assessed using balance then runs, reducing 2^l distinct sequences to not more than l^2 ordered equivalence classes of sequences which can be summarized in a single ordinal. These steps exorcize the curse of dimensionality.

Balance and runs are neither linear nor additively separable. For a concrete illustration, suppose the relevant sequence is 11011. The balanced action conditional on this sequence is 0, producing 110110. If the $t-1$ action had been 0, and the sequence 11010, the balanced action would be the same as before, producing 110100. Given older actions, the most recent action had no effect on the continuation.

If the model is intended as one of how people actually make decisions, the model needs to be simple. Balance and runs are intuitively transparent while suggesting specifications that conventional regressions may miss.

5.1. Comparison to learning theory

The most familiar alternatives to equilibrium behavior in mixed strategy games are learning theories. Two major modeling perspectives are reinforcement learning [14] and beliefs-based learning. A hybrid they call experience weighted attraction (EWA) is offered by Camerer and Ho [2]. They used M&S data to test EWA and Salmon [15] used simulations to assess the power of the M&S setup to distinguish between different kinds of learning behavior. Perhaps because the equilibrium action probabilities are so intuitively obvious, researchers have found little evidence of learning behavior in matching pennies.

In [9], O'Neill conducted an experiment with a repeated fixed-pair mixed strategy game having several actions, and one action had distinct payoff implications. O'Neill found little evidence of actions inconsistent with equilibrium, but Brown and Rosenthal [1] re-analyzed O'Neill's data and found actions inconsistent with equilibrium. Both Brown and Rosenthal and O'Neill [10] recognize that subjects have difficulty producing serially independent sequences, but neither propose a structural alternative to equilibrium.

A model of the first round of the O'Neill game motivated by framing and focal point issues comes from Crawford and Iriberri[18] while [5,12] are relevant theory papers.

It is hard to interpret the behavior identified here as learning because the identification strategy does not require evolution of behavior or convergence.

5.2. Contribution of this paper

Economists who research alternatives to equilibrium behavior in mixed strategy games may have sophisticated models, but they usually end up with straightforward tests. Most often they are

[18] V.P. Crawford and N. Iriberri, "Fatal attraction: focality, naivete, and sophistication in experimental hide-and-seek games", University of California San Diego, Economics Working Paper Series, 2004-12 (2004).

concerned with whether the cumulative action frequencies reflect the equilibrium action probabilities or best responses, or evolution toward either equilibrium or best responses.

When they consider sequences, economists most often explore a linear model by using dummy variables for each lag. I am aware of two types of tests which might better assess representativeness: a runs test and a kind of autocorrelation test. M&S report runs tests and have three rejections at the 5% level out of 20 tests, when one is expected. While Walker and Wooders [18] find stronger evidence for a runs test violation in their tennis data, Palacios-Huerta [11] finds none in his soccer data. Other researchers, including Croson and Sundali [3], have looked for excessive alternation in actions or responses to streaks.

Nash equilibrium reflects an exclusively defensive posture, an odd perspective on strategic interaction. I am aware of no previous study that distinguishes between defense and offense to analyze repeated fixed-pair mixed strategy games.

As far as I am aware, no economist has explored balance in a data set, though Rabin [12] considers it from a theoretical perspective. Representativeness and balance are familiar to psychologists, and I have borrowed from R&B at the conceptual level. R&B report a correlation statistic for an ordering close to the Rep decision rule, but propose no way of assessing statistical significance. The parametric bootstrap in [16] is the step that provides a test for statistical significance. Further, I define three new decision rules (BR Rep, C-Rep and BR C-Rep), develop tests for them and find evidence for two of them.

In most existing research, importance is measured in terms of a significant regression coefficient or correlation. In this paper, that is the first stage as reported in Table 7. Here significance is measured by exploitability. Exploitability is a step toward algorithms that might beat the alternatives in competition with human agents. This higher hurdle is not always cleared. In Table 11, Panel A, Column 2, for example, thousands of cases pass the first stage, but fail exploitability.

Acknowledgments

I acknowledge the suggestions of my advisor, Vincent Crawford, as well as helpful comments of numerous seminar participants and early readers, particularly Joel Sobel, Joel Watson, Mark Machina, Gary Charness and Philip Babcock. Anonymous referees significantly improved the paper.

References

[1] J.N. Brown, R.W. Rosenthal, Testing the minimax hypoethesis: a reexamination of O'Neill's game experiment, Econometrica 58 (1990) 1065.
[2] C. Camerer, T.H. Ho, Experience-weighted attraction learning in normal-form games, Econometrica 67 (1999) 827.
[3] R. Croson, J. Sundali, The gambler's fallacy and the hot hand: empirical data from casinos, J. Risk Uncertainty 30 (2005) 195.
[4] W. Edwards, Probability learning in 1000 trials, J. Exper. Psychol. 62 (1961) 385.
[5] D. Fudenberg, D.K. Levine, Conditional universal consistency, Games Econ. Behav. 29 (1999) 104.
[6] W.H. Greene, Econometric Analysis, fourth ed., Prentice-Hall, Englewood Cliffs, NJ, 2000.
[7] D. Kahneman, A. Tversky, Subjective probability: a judgment of representativeness, Cognitive Psychol. 3 (1972) 430.
[8] D. Mookherjee, B. Sopher, Learning behavior in an experimental matching pennies game, Games Econ. Behav. 7 (1994) 62.
[9] B. O'Neill, Nonmetric test of the minimax theory of two-person zerosum games, Proc. Nat. Acad. Sci. 84 (1987) 2106.
[10] B. O'Neill, Comments on Brown and Rosenthal's reexamination, Econometrica 59 (1991) 503.
[11] I. Palacios-Huerta, Professionals play minimax, Rev. Econ. Stud. 70 (2003) 395.

[12] M. Rabin, Inference by believers in the law of small numbers, Quart. J. Econ. 117 (2002) 775.
[13] A. Rapoport, D.V. Budescu, Generation of random series in two-person strictly competitive games, J. Exper. Psychol.: Gen. 121 (1992) 352.
[14] A.E. Roth, I. Erev, Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria, Amer. Econ. Rev. 88 (1998) 848.
[15] T.C. Salmon, An evaluation of econometric models of adaptive learning, Econometrica 69 (2001) 1597.
[16] S. Scroggin, Essays in dynamic uncertainty: behavioral economics, investment theory and law and economics, Ph.D. Thesis, University of California San Diego, 2005.
[17] D. Sonsino, Learning to learn, pattern recognition and Nash equilibrium, Games Econ. Behav. 18 (1997) 286.
[18] M. Walker, J. Wooders, Minimax play at Wimbledon, Amer. Econ. Rev. 91 (2001) 1521.
[19] J. Wooders, J.M. Shachat, On the irrelevance of risk attitudes in repeated two-outcome games, Games Econ. Behav. 34 (342) (2001).