Bounded Rationality in Randomization

Steven Scroggin*

This draft: November, 2003

Abstract

In repeated games with Nash equilibria in mixed strategies, players optimize by playing randomly. Players are boundedly rational in their randomization efforts. Arguably, they have no internal randomization facility and they fashion external randomization aids from the environment. By conditioning on past play, boundedly rational players exhibit a pattern. The pattern is characterized by cognitive limitations variously called local representativeness, the law of small numbers or the gambler's fallacy. I find one such pattern—balance then runs—in re-analysis of existing data for matching pennies experiments. While players and play are heterogeneous, the pattern makes prediction plausible. I implement prediction with a non-linear autoregression. Model 1 is a statistically and substantively significant tool for predicting behavior in matching pennies. There is evidence for two other behavioral models, both of which require some sort of sophistication—including a model of the opponent as boundedly rational.

Keywords: behavioral economics, matching pennies, local representativeness, mixed strategies, stochastic and dynamic games

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

All you know about luck for certain is that it's bound to change. — Bret Harte, "The Outcasts of Poker Flat," in *The Luck of Roaring Camp and Other Stories*

Economists believe that in repeated games with Nash equilibria in mixed strategies rational players should randomize. A matching pennies player should choose heads with 50% probability, and her choice should be independent of everything in her opponent's information set. Do players act according to the theory? Do they act randomly?

Here are two alternatives: (1) learning theory and (2) bounded rationality. While learning theory has its successes, Mookherjee & Sopher (1994)(M&S) found little learning behavior in the classic repeated fixed pairs matching pennies treatment. (The game and their experiment are described in section 3.) Subsequent learning theory tests on their data also explain little, Salmon (2002).

Does bounded rationality prevent players from randomizing? One form of bounded rationality in randomization, sometimes labelled the gambler's fallacy, generates negative autocorrelation. M&S found modest evidence of negative autocorrelation in own past play.

Another form of bounded rationality in randomization generally is called local representativeness. Local representativeness implies some conditions are perceived to be more random than others when in theory they are all equally likely. Rapoport & Budescu (R&B)(1992) make local representativeness operational by defining it in terms of a lexicographic weak ordering. The ordering is based on balance then runs (see section 1). Balance then runs may be picked up when testing other models, particularly the gambler's fallacy or learning theories. But balance then runs is non-linear over time, unlike alternatives.

I make three claims. Claim #1: The balance-then-runs local representativeness pattern ("representativeness") is statistically significant in the M&S data.

Accepting Shachat's finding (Shachat (2002)) that play is heterogeneous, I develop four models to predict behavior in the game. Each model has two steps: (1) steps one filters prior play for a pattern, and (2) step two generates a prediction. Claim #2: Model 1, representativeness, is a statistically and substantively significant predictor of play. Model 1 is a purely defensive strategy, a flawed attempt to play randomly.

Claim #3: Players exploit boundedly rational play in opponents. Counter representativeness is defined as playing the strategy counter to representativeness. Model 2, counter representativeness, is significant in R&B data. Model 4 best responds to players behaving as in Model 2. Model 4 is also significant in R&B data. Models 2 and 4 are different models of player sophistication. In both, players respond to some notion of their opponent's behavior. In Model 2 it is defense—an attempt to be inscrutable. In Model 4 it is offense—an attempt to outguess the opponent. Both are efforts to do better than Nash equilibrium would allow.

The next section describes representativeness, how to look for it, and the results in the M&S data. In section 2, predictive models are developed and tested in M&S and R&B data. The discussion ties this work to the existing literature; leads to thoughts about random variables and how we mimic them—and how we respond to others attempts; describes the curse of dimensionality in the information set and data mining; and offers extensions. Conclusions follow.

1 Representativeness in randomization: What it is and how to spot it.

The broad sweep of cognitive psychology is a trend away from finding and bemoaning irrationality, toward finding rational explanations for what appears on first glance to be irrational behavior, McKenzie (2002). One element of the older "heuristics and biases" literature is the "representativeness heuristic"—an error that arises from the difference between the character of populations and small random samples of them.

> Kahneman and Tversky (1972) asked participants which of two hospitals would have more days of delivering more than 60% boys. * * * Although the small hospital would be more likely to deliver more than 60% boys on a given day (due to sampling variation), participants tended to respond that the two hospitals were equally likely to do so. Kahneman and Tversky (1972) argued that representativeness accounted for the finding: Participants were assessing the similarity between the sample and the 50/50 generating process, which is equivalent for the two hospitals — McKenzie (2002).

The representativeness heuristic leads to a variety of cognitive errors. Local representativeness is one species of bounded rationality in randomization. Several overlapping concepts formalize local representativeness: negative autocorrelation, balance, runs, and anti-symmetry. I emphasize one form based on balance then runs, which implies some negative autocorrelation, but does not encompass anti-symmetry. (For more see section 3.)

In matching pennies truly random players generate every possible path with equal likelihood. The path HHHHH is as likely as any other 5-step path. A boundedly rational player may see different paths as having different likelihoods. A path with a proportion of heads and tails near 50/50 may be perceived to be more random than one with a more extreme proportion. So HTHTH, a path with 60% heads, may be perceived more random than HHHTH, with 80% heads. Conditional on the same proportions, paths with more runs maybe perceived more random. (For more on whether perceived randomness is monotonic in runs see section 3.3.) So HTHTH, having 60% heads and 5 runs, may be perceived more random than HHHTT, having 60% heads and 2 runs. Accordingly, having seen HHHTT, a player might think T next, generating THHHTT, (50 % heads) is more random than H next, generating HHHHTT (67 % heads). In sum, my model of representativeness has this intuition: Paths with a balance of heads and tails are perceived more random than unbalanced paths, and conditional on balance, paths with more runs are perceived more likely than those with fewer runs.

To make representativeness econometrically operational, more rigor is necessary.

An l-step *path* A_l is a sequence of l binary digits. For concreteness, digits are coded as "H" or "T." One example of a 5-step path is

$$A_5 = THTHH.$$

The fifth (oldest) step in A_5 is $a_5 = H$.

Let A_l be a path, A_l^H be the number of heads in path A_l, and A_l^T be the number of tails. Define the *balance* of a path,

$$\text{bal}(A_l) = \max(A_l^H, A_l^T).$$

The balance rule is the condition:

$$\text{bal}(A_l) > \text{bal}(B_l) \Rightarrow P[B_l] > P[A_l].$$

In words, if one path A_l has more of either outcome than path B_l has of either outcome, then B_l is perceived more random. For example $A_4 = HHHH$ is perceived less random than $B_4 = HTHT$ because the A_4 has 4 elements in its maximal category, and B_4 has only two. $P[\cdot]$ is a weak ordering implementing a subjective sense of probability; it need not be a probability measure.

Let $\text{runs}(A_l)$ be the number of runs in A_l, where a *run* is counted if $a_i \neq a_{i-1}, i \leq l$.

$$\text{runs}(A_l) = \sum_{i=1}^{l} I(a_i \neq a_{i-1})$$

where $I()$ is the indicator function—equal to one if true, and zero otherwise—and $a_1 \neq a_0$ by definition. The runs rule is

$$\text{run}[A_l] > \text{runs}[B_l] \Rightarrow P[A_l] > P[B_l].$$

Another example: $A_4 = HTHT$ is perceived more random than $B_4 = HHTT$ because A_4 has 4 runs, while B_4 has two.

A path A_l is more *representative* than B_l if it is more balanced or equally-balanced but has more runs.

$$\begin{aligned}\text{rep}(A_l) > \text{rep}(B_l) \quad \Leftrightarrow \quad & \text{bal}(A_l) < \text{bal}(B_l) \quad \text{or} \\ & \text{bal}(A_l) = \text{bal}(B_l) \quad \text{and} \quad \text{runs}(A_l) > \text{runs}(B_l).\end{aligned}$$

A representativeness player will perceive the representative path as more random and so play it. Then, over the course of many repetitions, the distribution of the player's empirical path frequencies will reflect representativeness: paths perceived to be more random will be played more often. Representativeness implies a weak ordering of path frequencies. "Weak" because some paths are equally representative. "Ordering" because representativeness orders path frequencies but does not measure them. Representativeness offers no quantitative insight into how frequent one path should be, either absolutely or relative to other paths: There is no cardinal information.

To test empirical path frequencies against the representativeness hypothesis, choose a rank correlation test. Two rank correlation statistics from non-parametric statistics apply, Kendall's τ and Spearman's ρ. Except as noted, all calculations are adjusted for ties. Sprent (1993) gives formulas and critical values.

The usual critical values depend on independent observations. Non-overlapping paths of play are independent too, but there are too few of them in the M&S data set to make a meaningful statistical test. Instead, I overlapped paths, and simulated critical values. Ten thousand sets of M&S data (800 random data points each) were created pseudo-randomly, paths extracted and critical values obtained from empirical quantiles of the rank correlation between frequencies of simulated paths and the representative weak ordering of paths. For example, for path length $l = 5$, 5% of the random data sets has Spearman's $\rho > 0.71465$. The critical value given independence is 0.294. The values differ because of dependence in overlapping paths, but may also be a feature of the specific finite sample, 20 observations

Table 1: Representativeness in M&S Data

Path Lengths l	Kendall's τ	Spearman's ρ
3	0.34	0.62
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%

of 40 rounds of play.

The alternative hypothesis is positive rank correlation, so I use a one-sided test. Row players win with matches; column players with mismatches. When considering column players, I reframe the column player data by labelling H as T and vice versa, so row and column players may be treated equivalently as match-seekers.

Rank correlations between empirical and representative path frequencies from the M&S data are reported in Table 1. There are no reference values for paths shorter than $l = 3$ because there are no more than four possible paths to rank when their lengths are so short. For long paths, the empirical path frequencies in the M&S data are arguably too small for a compelling diagnostic test: There are $2^9 = 512$ possible paths for length $l = 9$. Expected path frequency for path length $l = 9$ is about 1.2. It makes sense that if there is representativeness at length $l = x$ it would have an echo in $l = x + 1$ since the paths differ only in the $x + 1$'th element. So the eleven tests reported 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.

Having found evidence of a pattern, can we predict from it?

2 Prediction in Matching Pennies

2.1 Econometric Model

Let $R_{i,t}$ be the decision of row player i at time t, where $R_{i,t} \in \{-1, 1\}$. In the M&S data, as reframed, there are 20 pairs of players, and i uniquely identifies a pair of players. An opponent's choice is $C_{j,t} \in \{-1, 1\}, j \in \{1, \ldots 20\}, i \neq j$.

A general functional form for prediction of $R_{i,t}$ is

$$\begin{aligned} R_{i,t} &= f(R_{i,t-k}, C_{j,t-k}, X_i'\gamma, \epsilon_{i,t}), \\ & \quad i = 1, \ldots 20, j \in \{1, \ldots 20\}, j \neq i, k = 1, \ldots t-1, t = 1, \ldots T. \end{aligned}$$

Sets of variables $\{R_{i,t-k}\}$ and $\{C_{j,t-k}\}$ are R_i's information set. Player i may recall her prior play $\{R_{i,t-k}\}$, and her opponent's prior play, $\{C_{j,t-k}\}$. X_i is fixed effects, a vector of dummy variables, all set to zero except that the i'th is set to one. γ is a vector of coefficients on the dummy variables. $\epsilon_{i,t}$ is an error term.

The partitioning scheme is based on the idea that R_i will base her play on her own past play, or her opponent's, but not both at the same time. For more on the tradeoff between this strong assumption and tractability, see section 3.3.

First consider defense, in which R_i conditions only on her own past play:

$$\begin{aligned} R_{i,t} &= f(R_{i,t-k}, X_i'\gamma, \epsilon_{i,t}), \\ & \quad i = 1, \ldots 20, k = 1, \ldots t-1, t = 1, \ldots T. \end{aligned}$$

Since play is highly heterogeneous over players and play, apply a filter to sift predictable play from data for which we have no prediction. If we have no prediction, $\hat{R}_{i,t} = 0, (\hat{R}_{i,t} = 1/2$ for logit models). The filter takes the form of an indicator function:

$$I(\rho_{l,w} > x).$$

$I(\cdot)$ is one if the inequality holds, zero otherwise. The filter introduces a number of free parameters which could facilitate data mining. I have fixed the parameters, sometimes arbitrarily, to minimize this hazard. All are varied later as robustness checks.

The function $\rho(\cdot)$ is Spearman's ρ with tie adjustment. There are three other tests of rank correlation, $\hat{\rho}(\cdot)$, Spearman's ρ without tie adjustment, as well as $\tau(\cdot)$ and $\hat{\tau}(\cdot)$, Kendall's τ with and without tie adjustment. Spearman's ρ is the baseline statistic.

The rank correlation functions take two parameters, the path length l and the window w. Path lengths up to $l = 9$ are used. The test of representativeness is balance then runs, but for odd lengths l, there are no ties in balance, so representativeness is measured by balance alone.

The parameter w specifies the upper bound on how many paths are to be considered. For example, $w = 50$ specifies that the rank correlation between last fifty overlapping paths and the representativeness weak ordering is to be calculated. In the M&S data, with T=40, it is not obvious that w should be less than all the paths in the player's information set. The baseline value for w is $w = 50$. This choice calls for using all prior paths in the M&S data. This choice biases against useful filtering in the M&S data, hence against finding significance, when a player changes strategy.

I need a threshold number of paths on which to calculate rank correlations before predictions are practical. Let $\tilde{w}$ be this parameter. It is fixed at $l+1$, its theoretical minimum. This choice also biases against finding significance because correlations will be calculated even when there is only one path to compare to the representativeness distribution.

Next, x is a trigger value. If rank correlation exceeds x, conclude that behavior follows representativeness and predict the next choice. Set x at the rank correlation value that coincides with a 80 % probability of representativeness, where that 80 % value is from the simulated critical values derived in section 1. For path length $l = 5$ the 80 % value is $\rho = 0.422$. The filter applies to the player whose history is in the information set. Here it is own play; for Models 3 and 4 below, it changes.

Let $\hat{\text{rep}}_l(R_{i,t-1})$ be the representativeness prediction. Using R's immediately preceding path of length l, compare the representativeness of choosing 1 and of choosing -1 and predict the one which is more representative.

The specification becomes:

$$\begin{aligned} R_{i,t} &= f(I(\rho_{l,w} > x)\hat{\text{rep}}_l(R_{i,t-1})\beta, X_i'\gamma, \epsilon_{i,t}), \\ & \quad i = 1, \ldots 20, t = \hat{w}, \ldots T. \end{aligned}$$

$R_{i,t}$ is a limited dependent variable, suggesting use of a probit or logit functional form, among others. I use logit, with reference to OLS as a check.

$$\begin{aligned} R_{i,t} &= f_L(I(\rho_{l,w} > x)\hat{\text{rep}}_l(R_{i,t-1})\beta + X_i'\gamma + \epsilon_{i,t}), \\ & \quad i = 1, \ldots 20, t = \hat{w}, \ldots T, \end{aligned} \tag{1}$$

where

$$f_L(x) = \frac{e^x}{1 + e^x}.$$

The parameter of interest is β. β will be positive if the representative choice on filtered data is a significant predictor of $R_{i,t}$. Equation (1) is Model 1, representativeness.

Define counter representativeness as negative rank correlation, and predict counter-representative players by guessing that their choices will continue to be counter representative. Hence model 2:

$$\begin{aligned} R_{i,t} &= f_L\left(I(\rho_{l,w} < -x)(-\hat{\text{rep}}_l(R_{i,t-1}))\beta + X_i'\gamma + \epsilon_{i,t}\right), \\ & \quad i = 1, \dots 20, t = \tilde{w}, \dots T. \end{aligned} \tag{2}$$

Next, consider offensive partitions. A player may detect representativeness in her opponent—I assume she would do it exactly as I do—and best respond to it. This is Model 3:

$$\begin{aligned} R_{i,t} &= f_L\left(I(\rho_{l,w} > x)\hat{\text{rep}}_l(C_{j,t-1})\beta + X_i'\gamma + \epsilon_{i,t}\right), \\ & \quad i = 1, \dots 20, j \in \{1, \dots 20\}, i \neq j, t = \tilde{w}, \dots T. \end{aligned} \tag{3}$$

Finally, a player detecting counter representativeness in her opponent may best respond to that. This is Model 4:

$$\begin{aligned} R_{i,t} &= f_L\left(I(\rho_{l,t-l-1} < -x)(-\hat{\text{rep}}_l(C_{j,t-1}))\beta + X_i'\gamma + \epsilon_{i,t}\right), \\ & \quad i = 1, \dots 20, j \in \{1, \dots 20\}, i \neq j, t = \tilde{w}, \dots T. \end{aligned} \tag{4}$$

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

$$\begin{aligned} R_{i,t} &= f_L(I(\rho_{l,t-l-1} > x)\hat{\text{rep}}_l(R_{i,t-1})\beta_1 \\ &+ I(\rho_{l,t-l-1} < -x)(-\hat{\text{rep}}_l(R_{i,t-1}))\beta_2 \\ &+ I(\rho_{l,t-l-1} > x)\hat{\text{rep}}_l(C_{j,t-1})\beta_3 \\ &+ I(\rho_{l,t-l-1} < -x)(-\hat{\text{rep}}_l(C_{j,t-1}))\beta_4 \\ &+ X_i'\gamma + \epsilon_{i,t}), \qquad (5) \\ & \quad i = 1, \ldots 20, j \in \{1, \ldots 20\}, j \neq i, t = \tilde{w}, \ldots T. \end{aligned}$$

2.2 M&S Data

The baseline logit specification is defined as follows: Path length is $l = 5$. The filter omits observations for which the probability of representativeness is less than 80%, $\rho < 0.422$, calculated on all five-step paths up to time t. A constant is included. The results for Model 1, baseline specification are Table 2, Column 1. 115 observations spread among 8 of the 20 players passed the filter. In the absence of a consensus measure of goodness-of-fit for logit models, (*see,* Greene (2000), p. 831) I use conditional probability of successful prediction. It has a transparent intuition: 50% is worthless, 100% is perfect. The representativeness prediction was correct 37 times more than it was incorrect; correct 66.1% of the time. The z-score on $\hat{\beta}_1$ is 3.38, probability 0.0004. In this regression $\hat{\beta}_1 = 1.33$. This coefficient is not reported hereafter since it is its sign and statistical significance that matters.

Pair fixed effects are shown for the baseline model, Table 2, Column 3. The Model 1 β coefficient is slightly more significant, an F test for the fixed effects as a whole was not significant; for 2 players out of 20 fixed effects were

Table 2: M&S Data, Baseline, Portmanteau and Fixed Effects

Specification	(1)	(2)	(3)
Model	1	1,2,3,4	1
Path Length l	5	5	5
Filter	$\rho_{5,t} > 0.422$	$\rho_{5,t} > 0.422$	$\rho_{5,t} > 0.422$
	80%	80%	80%
$\tilde{w}$	6	6	6
Fixed Effects	No	No	Yes
Function	Logit	Logit	Logit
Results	(1)	(2)	(3)
Predictions	115	271	115
Net wins	37	53	37
Percent won	66.1	59.8	65.8
Pairs used	8	18	8
$\hat{\beta}_1$ z-score	3.38	3.42	3.57
Probability	0.0004	0.0003	0.0002
$\hat{\beta}_2$ z-score		0.94	
$\hat{\beta}_3$ z-score		0.46	
$\hat{\beta}_4$ z-score		1.64	

significant at the 90% level, consistent with insignificance.

The other models fare less well. The portmanteau model is characteristic, Table 2, Column 2. Models 2 and 3 are not statistically significant either here or separately. Model 4 is barely statistically significant both in the portmanteau model and when run alone at the 95% level, but the result is not robust to varying path length or the statistical measure.

Although there is little evidence of Models 3 and 4—players best responding to representativeness or counter representativeness—there is some evidence of players reacting to their opponents in a simpler way. Table 3, Column 5 shows Model 1, representativeness in own play with lags of opponent's play. No results are reported for success measures because I have no theory for how to combine representativeness with lags. Opponent's lags one and two are not significant alone (M&S have this result), but lags one and two are statistically significant with opposite signs when combined with Model 1, representativeness, in own play. The results are robust to dropping opponent's lags 3 and 4, which were not significant. Opponent's lags one and two remain significant at approximately the 95% level with variations in trigger level in the representativeness filter and variations in path length.

2.3 Robustness

This subsection is support for the Model 1 baseline specification, Table 2, Column 1. Though some specifications fit better in some respects, and others are worse, the baseline specification is the one on which I settle. There are three measures of the quality of predictions: (1) net successes, (2) probability of success given a prediction and (3) z-score.

Table 3: M&S Data, Baseline, Linear and Sophistication Effects

Specification	(1)	(4)	(5)
Model	1	1	1 + lags
Path Length l	5	5	5
Filter	$\rho_{5,t} > 0.422$	$\rho_{5,t} > 0.422$	$\rho_{5,t} > 0.422$
	80%	80%	80%
$\tilde{w}$	6	6	6
Fixed Effects	No	No	No
Function	Logit	Linear	Logit
Results	(1)	(4)	(5)
Predictions	115	115	
Net wins	37	37	
Percent won	66.1	66.1	
Pairs used	8	8	
$\hat{\beta}_1$ z-score	3.38		3.25
$\hat{\beta}_1$ t-stat		3.77	
Probability	0.0004	0.00008	0.0005
$\hat{C}_{t-1,j}$ z-score			1.66
$\hat{C}_{t-2,j}$ z-score			-1.52
$\hat{C}_{t-3,j}$ z-score			0.05
$\hat{C}_{t-4,j}$ z-score			0.11

Table 4: M&S Path Length l Variations

Panel A: Short Paths

Specification	(1)	(2)	(3)	(4)
Model	1	1	1	1
Path Length l	3	4	5	6
Filter	$\rho_{3,t} > 0.948$	$\rho_{4,t} > 0.491$	$\rho_{5,t} > 0.422$	$\rho_{6,t} > 0.290$
Results	(1)	(2)	(3)	(4)
Predictions	19	129	115	143
Net wins	11	29	37	29
Percent won	78.9	61.2	66.1	60.1
Pairs used		9	8	8
$\hat{\beta}_1$ z-score	2.33	2.52	3.38	2.41
Probability	0.01	0.006	0.0003	0.008

Panel B: Long Path Lengths

Specification	(5)	(6)	(7)
Model	1	1	1
Path Length l	7	8	9
Filter	$\rho_{7,t} > 0.223$	$\rho_{8,t} > 0.156$	$\rho_{9,t} > 0.121$
Results	(5)	(6)	(7)
Predictions	143	137	116
Net wins	39	25	24
Percent won	63.6	59.1	60.3
Pairs used	8	9	8
$\hat{\beta}_1$ z-score	3.22	2.12	2.22
Probability	0.0006	0.017	0.013

Consider varying the path length, with trigger level set at 80%. This is summarized in Table 4. The results are statistically significant for all paths sizes from $l = 3$ to $l = 9$. Net successes range from 11 to 39. Success rates vary from 59.1% to 78.9%. Odd length paths do better. The first and second best on all three of my ranking measures are odd-length paths. On the other hand, path length $l = 2$ was not statistically significant ($\tilde{w}$ was set to 4 to get it to run).

Consider varying the trigger level x. Net successes vary from 14 (no trigger) to 43 ($x = 0.20$) and then fall to 16 ($x = 0.80$). The probability of a success given a prediction is only 51% with no filter, but increases smoothly to 77% for $x = 0.80$. The z-score varies from 0.54 (no trigger) to 3.87 ($x = 0.50$) and then falls to 2.76 ($x = 0.80$). In this dimension, one can obtain non-significant results for $x < 0.20$. This is evidence that a filter is necessary.

Consider varying the rank correlation statistic. There are four to choose from. In the baseline specification they are all significant; Spearman with no tie adjustment is worst, with z-score 2.90. Win rates range from 60-75%; net wins 21-41. No one measure dominates.

Similar robustness tests were conducted on increasing $\hat{w}$ from its minimum value or decreasing w so that it was significantly constraining. Increasing $\hat{w}$ does not matter unless it begins to decrease the amount of available data significantly. In the baseline specification with $\hat{w} = 15$ the z-score for the representativeness variable was 3.20. Decreasing w dramatically also has little effect. For the baseline specification with $w = 9$, the z-score is 3.34. For even smaller windows, the results begin to deteriorate.

The favorable results for Model 1, representativeness, are robust. Any path length anywhere near $l = 5$ works. Trigger values for x well below and above the baseline level work. Any of four measures of rank correlation work. Varying the beginning point $\hat{w}$ and the window w make little difference. Fixed effects were unimportant—perhaps the filter successfully supplants them. OLS results are about the same are logit. Model 1 stands out relative to the other models. The number of its predictions and z-scores were consistently larger than that of other models.

Also of note: The path length l plus about 10 rounds is sufficient to generate optimal predictions. And the best fits are for $l = 5$ or $l = 7$. So Model 1 does not require lots of data. This indirectly suggests that Rabin's N (see discussion) is roughly 16.

2.4 Simulation

As another check on hidden dependencies due to overlapping paths, for example, or data mining, I simulated a random data set and ran in through paces similar to those for the M&S data. I tried 10 regressions with different path lengths, triggers, statistical measures, and windows, all for Model 1, representativeness. The regression with the best fit had probability 11%, though the sign was wrong.

2.5 R&B Data

After completing work on the M&S data, I turned to the R&B data. R&B is another matching pennies experiment conducted in a different country with the experimental protocols of another discipline: psychology. Details are

Table 5: R&B Data, Path Length $l = 5$

Specification	(1)	(2)	(3)	(4)
Model	1	2	3	4
Path Length l	5	5	5	5
Filter	$\rho_{5,50} > 0.422$	$\rho_{5,50} > 0.422$	$\rho_{5,50} > 0.422$	$\rho_{5,50} > 0.422$
Results	(1)	(2)	(3)	(4)
Predictions	3558	1158	3630	1129
Net wins	568	54	-84	85
Percent won	58.0	52.3	48.8	53.8
$\hat{\beta}$ z-score	9.46	1.56	-1.44	2.46
Probability	0.0000000	0.058	0.07	0.007

given in section 3. The same baseline parameters should work, except for the trigger value x, whose interpretation is specific to the M&S sample size and number of rounds. Nevertheless, for consistency, I retained the same trigger values.

Results for the Model 1 baseline specification are given in Table 5, Column 1. This is a larger data set; 3,558 predictions were made with 568 net wins (58%) and z-score 9.46, probability 0.0000000. Models 2, 3 and 4 are the next three columns. Model 2, counter representativeness, just misses statistical significance in the baseline model. Model 3, best response to representativeness, is not significant and has the wrong sign. Model 4, best response to counter representativeness is statistically significant, but not nearly so salient as Model 1.

The R&B game is almost four times longer than the M&S game (for those who went the distance). Perhaps in a longer game, longer paths matter. And a longer game with more pairs supplies more data. Longer paths means more possible paths, and allowing the path space to ramify is more promising in a larger data set. Table 6, Panel A is the baseline specification for the four models, but with path length $l = 7$. Model 1 has equally strong results. However, the borderline results for Model 2 are improved. With length $l = 7$, Model 2 counter representativeness is less prominent than Model 1 representativeness, but it is clearly present. Model 3, best response to representativeness, retains the wrong sign and remains insignificant. Model 4, best response to counter representativeness remains significant.

The three significant models are combined in Table 6, Panel B, Column 1. Little changes, suggesting the models are largely independent of one another. Column 2 adds the first lag of opponent's play. The first lag is significant and reduces but does not eliminate the significance of Model 4. Longer lags were insignificant.

Table 7 repeats the exercise for path length $l = 9$. Model 1 is as before, although it makes about 20% more predictions and gets them right with almost the same probability as with $l = 5$. Model 2 counter representativeness is now also extremely significant. A subroutine in Gauss, QNewton Version 5.0.14 could not solve Model 3 because a matrix was complex—perhaps consistent with its being insignificant. Model 4 is weaker than with $l = 7$, though still significant.

2.6 Summary of Matching Pennies Results.

Some play is consistent with the null hypothesis of randomness; however, some play is predictable. Model 1 was able to pick up on patterns and forecast from them. In many R&B specifications, when it guessed, it guessed

Table 6: R&B Data, Path Length $l = 7$

Panel A:

Specification	(1)	(2)	(3)	(4)
Model	1	2	3	4
Path Length l	7	7	7	7
Filter	$\rho_{7,50} > 0.223$	$\rho_{7,50} > 0.223$	$\rho_{7,50} > 0.223$	$\rho_{7,50} > 0.223$
Results	(1)	(2)	(3)	(4)
Predictions	4018	1194	4070	1138
Net wins	632	142	-60	68
Percent won	57.9	55.9	49.3	53.0
$\hat{\beta}$ z-score	9.92	3.98	-0.95	2.01
Probability	0.0000000	0.00003	0.17	0.02

Panel B

Specification	(1)	(2)
Model	1,2 4	1,2,4,lag
Path Length l	7	7
Filter	$\rho_{7,50} > 0.223$	$\rho_{7,50} > 0.223$
Results	(1)	(2)
Predictions	5576	
Net wins	806	
Percent won	57.2	
$\hat{\beta}_1$ z-score	9.97	9.96
$\hat{\beta}_2$ z-score	4.00	3.97
$\hat{\beta}_4$ z-score	2.21	1.83
$\hat{C}_{t-1,j}$ z-score		-3.11

Table 7: R&B Data, Path Length $l = 9$

Specification	(1)	(2)	(3)
Model	1	2	4
Path Length l	9	9	9
Filter	$\rho_{9,50} > 0.121$	$\rho_{9,50} > 0.121$	$\rho_{9,50} > 0.121$
Results	(1)	(2)	(3)
Predictions	4210	1147	1067
Net wins	622	225	67
Percent won	57.4	59.8	53.1
$\hat{\beta}$ z-score	9.54	6.48	2.02
Probability	0.0000000	0.0000000	0.02

correctly four times in seven on average, over thousands of rounds.

Model 2, based on counter representativeness, was successful in the larger R&B data set, but not in the M&S data set. A player might apply Model 1 to her opponent's play, detect representativeness and best respond to exploit it, thereby becoming predictable in turn. This is Model 3. Model 3 was an unexpected failure. Model 4, a best response to opponents using Model 2, was an unexpected success in the R&B data set given failure of Model 3. Overall, the strongest results obtained were for path length $l = 7$ in the R&B data, Table 6 and the baseline specification path length $l = 5$ in the M&S data, Table 2, Column (1).

3 Discussion

3.1 Links to the Literature

In matching pennies Row and Column choose a binary variable at the same time; for concreteness, suppose they choose heads or tails. If they match—

both heads or both tails—then Row wins one unit and Column wins nothing. If they mismatch, then Column wins one unit and Row wins nothing. That's one round of the stage game, Figure 1. In experiments there are numerous stages between the same players.

The school yard version of matching pennies is a zero sum game. In the economics experiments the game is structured as a positive sum game. This structure avoids the possible impact of different preferences for gains and losses. Players are usually given an endowment, and it also avoids the risk that one player might lose his entire endowment. On the other hand, psychologists feel such possibilities add spice to the mix, Rapoport & Budescu (1992). In R&B 45 pairs of undergraduates from the University of Haifa played 150 rounds with an initial endowment of 20 New Israeli Shekels. Some lost their endowment and the game terminated in less than 150 rounds. R&B also experimented with two other conditions, which are not important here because they involved no strategic interaction. M&S players (Moohkerjee & Sopher (1994)) were masters students in economics at the University of Delhi. Ten pairs of players played 40 rounds of two treatments. In treatment 1 players did not know the payoff matrix. Only the second treatment is of interest here. Treatment two was the standard matching pennies game. The payoff matrix was common knowledge and payoffs were announced after each round. I thank both Barry Sopher and David Budescu for their data.

Nash equilibrium captures rational behavior in matching pennies. Nash equilibrium in this game is intuitively easy: Row plays heads with i.i.d. Bernoulli probability 1/2. Column does the same. If play is not independent, then the other player should be able to exploit the information contained in

	H	T
H	0 1	1 0
T	1 0	0 1

Figure 1: Matching Pennies Normal Form Game

past play. If it is not Bernoulli probability 1/2, then the other player should be able to exploit the difference by always best responding to the more likely outcome. M&S apply a number of econometric tests in cross sections and time series to show that behavior in their experiment is largely consistent with randomness and Nash equilibrium behavior.

The core idea of learning theory is that players experiment with different strategies, observe the results and then modify their behavior, Roth and Erev (1998). Fudenberg & Levine (1998) is a recent text. Reinforcement learning is a stimulus/response model; behavior learning entails responses to a model of opponent behavior. The M&S treatments were designed to distinguish reinforcement from behavior strategies experimentally, but the design may have been contaminated by learning about the game in Treatment 1. Compare Oechssler & Schipper (2003). Beyond noting that play in their treatments differ, M&S do not find learning behavior. M&S data have been reviewed by others. Camerer and Ho (1999) used it to test their EWA parametric learning theory. Salmon (2001) used it to test the power of learning theory econometrics to distinguish between different kinds of learning behavior. Despite this effort, the evidence for learning in M&S treatment 2 is weak.

O'Neill (1987) conducted a slightly more complicated game experiment. There were several alternatives, and one had larger payoff implications than the others. O'Neill found little evidence of play inconsistent with Nash equilibrium. Brown & Rosenthal (1990) re-analyzed O'Neill's data and found play strikingly inconsistent with randomization. O'Neill (1991) criticized their approach saying that all theories are precisely wrong and their tests did not reject randomization in favor of an alternative. I claim that representativeness is a successful alternative. Further, matching pennies may be the stiffest challenge one can place before a bounded rationality theory because play in matching pennies is so obvious. If players can randomize anywhere, they can randomize in matching pennies. Since they do not consistently randomize in matching pennies, bounded rationality in randomization must be a common phenomenon in untrained subjects lacking a randomization device.

A variety of two-person zero sum games were studied in Binmore, et al. (2001). The pairings were reshuffled after each game, a design that obscures sequential behavior in favor of other phenomena.

Bounded rationality in randomization or local representativeness is not entirely new to social science, Camerer (1995). Predictability requires dependence between past and future play. It has long been recognized that some sort of dependence exists, Edwards (1961), Brown & Rosenthal (1990). The negative autocorrelation variety of local representativeness is easy to test for and has been found in many studies, Bar-Hillel & Wagenaar (1991). Local representativeness has also been characterized in terms of balance, runs and anti-symmetry tests. The test of local representativeness I consider is inspired by Rapoport & Budescu (1992),(1997). In their work they go so far as to calculate one of the rank correlation coefficients I use against a weak ordering closely related to the one I derive from their work. Lacking critical values for overlapping paths, they do not perform a formal test, nor do they offer a predictive algorithm.

Balance conditional on a fixed path length is deterministic and eliminates some paths completely. This is fine for a forecast, but it is dubious as a model of actual behavior since some paths are seen in data that are not representative conditional on any possible history. One might simply add an error term, but if the point is to understand how players generate noise, is it appropriate to just assume some noise?

If short-term memory is itself stochastic, then players' behavior may be representative conditional on a stochastic memory constraint. This suggests a question about where and how the random memory constraint arises, but at least it pushes the problem a step back. In simulations, R&B show that a stochastic memory constraint avoids eliminating some paths. They provide evidence that this weak ordering is robust to the length of short term memory, and its stochastic structure. If the empirical path frequency ordering is close to the representativeness ordering, then, they argue, representativeness is confirmed. Even if this is wrong, if the weak ordering shows up consistently, there is something in the data that requires explanation.

The balance test aggregates history non-linearly. Suppose, for illustration, a player uses representativeness in 5-step paths, and the path is $HTTTH$. Continuing with H would be perceived more random because the conditioning path is 40% heads. The most recent choice H had zero impact on this decision, conditional on more remote history. On the other hand, if the path were $HHTTH$, the most recent choice H determined the next choice, T. Balance then runs to resolve ties is still more non-linear. Except for conditioning paths of length one, balance then runs is quite distinct from autocorrelation.

Whether it takes the form of simple autocorrelation or the linear combinations that underlie learning theory, there is no reason to expect players to be as mesmerized by linearity as researchers are. Linear regressions require a lot of data and computation. Simple, but non-linear, alternatives that may do about as well. The "take the best" heuristic championed by Gigerenzer, *et. al* (1999) ("Gigerenzer") is an example. Gigerenzer claims that boundedly rational players may use a variety of simple rules, ranked by importance. In performing a task, they apply the rules one at a time, and the first to provide a definite answer is the decision. He shows that, at least in some limited circumstances, the outcome is as good or better than the most information-intensive and sophisticated analysis. Such rules are candidate hypotheses of how people actually make decisions. Balance then runs is a "take the best" rule.

Applying the balance then runs version of local representativeness, or some other model, and finding that player strategies are indeed dependent on history, is only the first step. A more ambitious question: Can play be predicted sufficiently to play "offense," and thereby win more often (in expectation) than with defensive randomness? A theory that purports to explain behavior should suggest an algorithm that does better than the Nash equilibrium against the distribution of behaviors of real players. Shachat and Swarthout (2002) represents an initial attempt to build such an algorithm. Model 3 is another. An experimenter playing this model against subjects might do well.

Rabin (2002) offers an appealing variation of the balance test. He suggests players view data generating processes as draws from a urn with N coins (call them heads and tails for our purposes). If the urn started out with equal proportions and history is $HTTTH$, then the urn now has more heads than tails, so heads is perceived more likely—consistent with the balance test. Knowing N, one can calculate probabilities and so obtain cardinal information about perceived randomness. In a fully-developed model N may be a parameter found econometrically. N may be of intrinsic interest as a measure of the magnitude of boundedness in rationality. What is its mean? How does it vary over players, training, time and randomization problems? Rabin, however, gives up something with this theoretical advance: Rabin's model of how we make and assess noise has noise buried within it as a primitive. Based on N and history, he delivers a probability. Model 1 is deterministic.

Rabin's models are significantly more general than the class of models here. The implicit assumption here is that θ, the cross-sectional probability of H, is known and $\theta = 0.5$. One might allow this probability to be an unknown random variable. Bayesian updating of θ is already in the literature, Shachat & Swarthout (2003). Rabin allows for boundedly rational updating where the bound arises from Row's finite urn.

In this context, what is noise, and how might we go about making it?

3.2 Random Variables and the Taxonomy of Types

> O, what a tangled web we weave when first we practice to deceive!
>
> — Sir Walter Scott

Late bloomers in mathematics, random variables are a bit mysterious. Consider this sequence:

$$26535897932384626 43\ldots$$

What is the next number in the series? The properties of this series have been studied; the distribution of the digits is uniform, there is no standard form of serial correlation between the digits. If you do not know the trick it is a random variable. But if you know the trick, its easy: These are the 6'th through 25'th digits of π, and the next digit is 3. Much of what we are tempted to treat as random is deterministic, if only we knew the trick.

Much of what is not predictable in practice is thought to be chaotic. The evolution of chaotic systems is deterministic but unknown because parameters are not precisely known, and ultimately—courtesy of the Heisenberg Uncertainty principle—can not be known with sufficient precision. We may model chaotic systems with random variables for lack of an alternative. Aside from quanta, are there any true random variables?

Can you make random draws by introspection, out of the unconscious void? I can not. If there are no true random variables, it is no surprise we do not carry a random number generator around with us in our heads. To create apparent randomness I suggest we use external aids. The physics of coin tosses are difficult enough that the outcomes, while perhaps deterministic, are random for our purposes. A computer's random number generator is deterministic, but obscure enough to seem unpredictable.

Consider a subject playing matching pennies defense, aiming at inscrutability. Unless she's Shachat's subject (Shachat 2002), there is no obvious external randomization aid, and precious little serendipity at hand from which to fashion it. Accordingly, she may have recourse to the only stuff around, the history of her own play and that of her opponent. In the spirit of modern cognitive psychology, she is doing the best she can in a claustrophobic environment, generating a series from the information available—recent play. Given some short-term memory constraint and the sparse environment, what else could one do? If the answer to this rhetorical question is "not much," then perhaps this species of bounded rationality is not so "bounded" after all.

There need not be just one strategy, and all of them can happen in the same experiment. Player heterogeneity is consistent with other work. When uncertainty takes the form of the behavior of another player, players may be classified by types, Costa-Gomes, Crawford and Broseta (2001). Most players can be described as taking their opponent as assuming a diffuse prior over their behavior, or (one step deeper) best responding to a diffuse prior in the player's behavior. A few iterate the process further, toward Nash equilibrium, but there is evidence that it is not optimal outthink your opponent too far. The most sophisticated type will take into account the distribution of types among opponents and best respond to it.

Results specific to mixed strategies (Shachat (2002)) suggest the play of players also varies. Play is sometimes deterministic, though the deterministic choice may appear random to the opponent. At other times players use a randomization device if one is at hand. Play is non-stationary. There is no evidence that non-stationarity reflects learning behavior. Perhaps the source is changes in strategy which are intended to be a form of meta-randomness. The player has a menu of strategies, none of which are truly random or even very complex, but if she frequently switches among them in a complicated fashion the result may be very difficult to decipher.

When, with Bret Harte, the only thing you are sure of is that your luck will change, it is natural to look for snippets of history that display patterns rather than to look player by player, much less the whole data set. The filter used by the predictive algorithms respects player heterogeneity in two ways: (1) it allows players to differ from one another, (2) it allows players to vary over time. Some play exhibits the representativeness pattern (Model 1). Some shows sophisticated defense (Model 2) or offense (Model 4).

Model 2 counter representativeness is a defensive alternative to representativeness. Counter representativeness requires excessively persistent play. I view this as a strategic response to thinking of randomness using something like a balance test. If a balanced path is perceived random, a Model 2 player plays counter to that to (1) be "more random" than random or (2) to fake out the opponent. Either way, it is a defensive strategy, but a sophisticated one in that it is conscious of an opponent's presumed approach to the problem.

On the other hand, an offensive player is shuffling through history hunting for a hint. She seeks patterns in seeming noise. Knowing that its source is human is encouraging, perhaps the pattern is simple enough unearth. But

> [t]here is no systematic way to get it. One person could look at the pile of square wave tracings and see nothing but noise.

> Another might find a source of fascination there, an irrational feeling impossible to explain to anyone who did not share it. Some deep part of the mind, adept at noticing patterns (or the existence of a pattern) would stir awake and frantically signal the quotidian part of the brain to *keep looking* at the pile of graph paper. The signal is dim and not always heeded, but it would instruct the recipient to stand there for days if necessary, shuffling through the pile of graphs like an autist, spreading them out over a large floor, stacking them in piles according to some inscrutable system, pencilling numbers, and letters from dead alphabets, into the corners, cross-referencing them, finding patterns, cross-checking them against others * * * (Neil Stephenson, *Cryptonomicon*, p. 117.)

until despair, psychosis or epiphany breaks the deadlock.

Suppose, like Churchill's code breakers during World War II, she finds a pattern. The next stage in sophistication is just to use it and assume your opponent is not hunting for hints too. This is very much like Models 3 and 4. If Model 1 representativeness is common, and it is, one would expect opponents to pick up on it and exploit it. However, the failure of Model 3 is explicable if players cannot differentiate an opponent's successful randomization from an opponent's representativeness. If representativeness is a sound strategy, it must be hard to distinguish from noise.

On the other hand, counter representativeness Model 2 is playing one strategy with higher probability than another, at least locally, and Model 4 is best responding to Model 2. The significance of Model 4 was undermined when the opponent's lagged plays were added separately—and the lags were sometimes conditionally significant. Perhaps Model 4 is a mis-specification of a model that best responds to recent non-optimal probabilities in the opponent. If Model 4 is more successful than Model 2, and it sometimes is, it may mean that opponents perceive more persistence in persistent players than actually exists.

Simple offense as in models 3 and 4 suppose that the opponent you exploit will not see the pattern you create in exploiting his pattern. This may not be a winning strategy.

> We know everything * * * . We receive Hitler's personal communications to his theater commanders, frequently before the commanders do! This knowledge is obviously a powerful tool. But just as obviously, it cannot help us win the war unless we allow it to change our actions. * * * [A]t which [point] information begins to flow from us back to the Germans (*ibid*, p. 124).

Consider Churchill at the bombing of Coventry—when he choose not to use information that could have saved a city lest he betray his sources and lose access to still more vital information to come. What would a Churchill look like? That would require a player (a "German") whose opponent (Churchill) had a model that found and responded to predictability—like Model 4. Churchill's offensive pattern would be detectible by the researcher, but at some cost kept too subtle for his German counterpart, and that subtlety too would be detectible. (The quotes above are from a work of fiction, and some of those who deal in fact think the Coventry story belongs to fiction; consider it a metaphor.)

I four four kinds of play: random and Models 1, 2 and 4. I found neither Churchill, nor a metaphorical German scrutinizing him. But it is early. The experiments analyzed here were not designed with this is mind, nor are the models subtle enough to capture such guile.

3.3 Method: The Curse of Dimensionality and Data Mining

The patterns of play found here are non-linear. Simply grinding through regressions is unlikely to unearth them. Designing an econometric strategy requires a fresh look at the statistical problem.

It is obvious. Matching pennies is a very simple game. It has only four possible outcomes per round. Predict play in the 40'th round based on the previous 39. There are 4^{39} possible histories of play. If players randomize then all histories are equally likely and equally useless as predictors for round 40. For a completely non-parametric test, treat each possible history as a distinct category—4^{39} dummy variables in, say, a logit model. The number of 40-round games necessary to test this model has a magnitude somewhere between the U.S. debt in dollars and the universe in atoms.

The psychology literature suggests that players retain at most 7 distinct items in sort-term memory, see Rapoport & Budescu (1997). Which leads to the question, what counts as an item? Perhaps players retain running totals or some other aggregates of a round; or perhaps an item is as petite as what the opponent played last time. For the sake of discussion, one round of play is an item, and short-term memory is seven rounds. Recall that baseline specification is $l = 5$. Five is a nice number in terms of the cognitive model of limited short-term memory, yet long enough to be quite distinct from a linear model. But seven is almost as good in M&S and better in R&B. Make it seven.

Seven rounds yield $4^7 = 16,384$ paths through the tree of possibly-remembered play. It is better than 4^{39}, but still far to many. In a game of 40 rounds, only 33 seven-step paths will be realized. With 10 pairs, as in M&S, at most 330 distinct seven-step paths exist in the data. More than 16,000 of the possible paths will not occur at all, which makes calculation of their relative frequencies awkward.

While there are 16,384 relevant paths, it is plausible that players treat some of them as functionally identical. An intuitively appealing grouping is to look at offense versus defense. Suppose Row seeks to befuddle Column's efforts to guess Row's model. Playing defense, Row uses a partition of path space conditioning on her own past play only. Learning theory makes a similar division between reinforcement learning and behavior strategies. A reinforcement strategy conditions on own past play and payoffs. On the other hand, if Row plays offense, her behavior will induce a partition of the path space conditioning on Column's play only. The goal of offense is for Row to win by using past Column play and a model of his behavior to forecast him. Offense has an analogy in behavior strategies. Offensive and defensive partitions preclude decision-making based, for example, on Row play for the first 3 lags, and Column play for the next 4. But the gains in tractability are large. Now instead of 16,384 paths there are $2^7 = 128$ paths of offense and a like number of defense. Yet 128 dummy variables is still too many.

Assume that the effect on choice in period t of period $t-n$ is independent of the effect of period $t - m, m \neq n$. Then a linear specification can help. Seven dummy variables will do, one at each period for seven lagged periods. Every path is a linear combination of these seven choices. This structure is a reasonable and parsimonious response to the problem of dimensionality. Learning theories and M&S's tests are all based on linear combinations of past play. On the other hand, true randomness is (trivially) linear, and the dummy variable specification is linear as well. If behavior is non-linear, by assuming linearity we assume a feature of the null in order to test it—which may make it harder to disprove.

Consider other partitions of the path space. Then: (1) one can choose a partition that embodies a theory of actual decision-making, like representativeness or Rabin's model and (2) one can choose a partition that reflects a more thoughtful trade-off between the curse of dimensionality and loss of information from grouping paths.

Balance partitions paths into equivalence classes and then ranks the classes, like preference orderings partition consumption space. Consumption bundles are separated into equivalence classes, the indifference curves of consumption theory. Preferences impose an ordinal ranking among bundles on different indifference curves. Analogously, balance divides paths into equivalence classes of equal perceived randomness, and supplies an ordinal ranking of the classes. In seven step defense 128 paths fall into 8 equivalence classes, and these classes are ranked from 1 to 8 in perceived implication for the next round. Now we are in business. This is more parsimonious than dummy variables, yet incorporates a non-linear specification. What is more, since it is supposed to reflect bounded rationality, the model ought to boil down to something simple. These equivalence classes are used to diagnose patterns, but not to predict. Prediction is easier: Make the more representative choice.

Balance then runs is but one model for boundedly rational randomization. One might use balance alone, considering only paths of odd length so there are no ties, or adding some other tie-breaker. Another alternative: the most representative long path is simple alternation, $HTHTHTHTHT\ldots$. Perhaps few would think that path particularly random, and a better test would be based on how close the number of runs is to its expected value, conditional on randomness (or some boundedly rational variant). Yet another way: Take the basic unit as more than one round. If the basic unit is two rounds then the strategies are HH, HT, TH and TT; apply a representativeness model to these. But this may be too clever by half. My choice for this paper—balance then runs—has the benefit of a finer weak ordering than balance alone, while being simple relative to alternatives.

Now turn to the plausibility of the results. The probability that the baseline specification, Table 2, Column 1 is a chance occurrence is less than four in ten thousand, according to the z-score. In probing the quality of the phenomenon here, I confess to running many regressions, but not quite that many. And the diagnostic tests in section 1 are not subject to this concern.

Still, data mining concerns are serious. Here is a list of the elements of the work that could have been manipulated in search of a result: path length l, trigger level, x, correlation statistic, ρ, minimum observations value, $\hat{w}$, window size, w, fixed effects (yes/no), logit v. OLS, and permutations of Models 1 through 4. However, I have reported the effects of varying all these parameters. Without exception, the Model 1 results are robust to modest changes in parameter values and inconsistent with simply trying everything until something worked. One might also manipulate the data directly, omitting troublesome players or time periods. Here all the data were submitted to the filter. Or one might use a different path length l for the filter than for the representativeness prediction. I did not.

The Model 1 baseline specification is so robust it might be too good to be true. Other models, by their contrast, provide a reality check. Models 3 *never* worked, even at a modest level of statistical significance. Models 2 and 4 were relatively robust in the R&B data set, but not M&S.

The R&B data is also a robustness check—to see whether the baseline result was a fluke, but also to test whether the M&S data set was a fluke. It passed. The z-scores for R&B are even higher than in the M&S data. Finally, I simulated pseudo-random data sets and tried to mine them. If the analysis is valid, this mining project should fail. It did.

3.4 Is it Economics?

If psychology is the study of what people *do*, economics is the study of what they *ought* to do. This paper straddles the fence. Representativeness is a bounded rationality model—what people do rather than what they ought to do. Yet representativeness is also rational optimization constrained by absence of an internal randomization device, and limited memory. While the constraint is not in the stars—it is in ourselves—it remains exogenous. Such constrained optimization is the root of economics. An alternative economic interpretation: Optimize when an opponent is boundedly rational. Models 3 and 4 are based on this approach.

3.5 Extensions

The data do not admit a test of how players play against a true pseudo-random opponent. The data reviewed here involve human players facing other human players, who may or may not be completely unpredictable. An experiment with human players playing against a computer playing pseudo-random strategies might provide insight. Shachat & Swarthout (2003) have such an experiment, but their analysis heads in a different direction.

My analysis could not determine whether predictability persists in the presence of an opponent taking advantage of it. It is not clear whether players are aware of representativeness at all, or whether the patterns present "hunches" or "intuitions" whose source is obscure to the player. The evidence of sophisticated behavior in this data suggests that players may not respond to Model 3 differently than they respond to a typical human player or to randomness. Conversely, there is experimental evidence that players can learn to avoid autocorrelation, Camerer (1995). For a direct test, try an experiment in which human players confront Model 3: it plays pseudo-randomly when representativeness is not significant and otherwise best responds.

It might also be interesting to have the computer play representativeness, or representativeness conditional on stochastic path length, to see whether players can detect and exploit this predictability.

Shachat's game (Shachat (2002)) has more stage game structure than matching pennies, but it would be interesting to supply a matching pennies version of his 'shoe' and see if representativeness is associated with how the shoe is used. I am concerned, however, that providing the shoe and instruction on how to use it will affect behavior. It could be part of learning how to be random.

How do pairs differ? Is representativeness in a player correlated with representativeness in the opponent? Probably not. That Row is representative has no implications for Column. This fits with the surmise that opponents have difficulty separating representativeness from random play.

The Nash equilibrium strategy is a random walk. It has a unit root. Does experimental data have a unit root. Is this another source of predictability?

There are situations other than matching pennies in which humans have an incentive to generate randomness. One unanswered question (among many) is whether market forces punish representativeness severely enough to expunge it from the market, Camerer(1995). In the absence of a structural model, my hypothesis is simple: If the market exhibits short-term predictability, an arbitrageur will exploit it, if they can profitably. Expect representativeness to be more significant where transactions costs relative to potential rewards are larger—in higher frequency data for smaller issuers.

The filter might be modelled in light of Rabin (2002). Each equivalence class of paths may have a well-identified distinct probability for, say, H in the next stage. When N is small and the equivalence class extremely unbalanced, the probability of successful prediction should will be large. If N is found to be stable, predictability could be predicated on N and a single path rather than the computationally intensive filter used here. The model might also be extended to allow for boundedly rational updating of θ.

4 Conclusions

One who sees glasses as half full will conclude that about half of us are really good randomizers. Others, while predictable, use a non-linear technique that is not half bad. While I have spotted their trick, few of their opponents did. One who sees glasses as half empty will see that some of us can be gamed, and perhaps even wonder how such boundedly rational genes made it this far.

Defensive representativeness is a prominent phenomenon in matching pennies; other patterns were statistically significant but not as striking. So there are a few answers here and many questions. This paper is more akin to the first word than the last.

References

[1] Bar-Hillel, Maya and Willem A. Wagenaar, "The Perception of Randomness," 12 *Advances in Applied Mathematics* 428 (1991).

[2] Binmore, Ken, Joe Swierzbinski, and Chris Proulx, "Does Minimax Work? An Experimental Study," 111 *Economic J.* 445 (2001).

[3] Brown, James N. and Robert W. Rosenthal, "Testing the Minimax Hypothesis: A Reexamination of O'Neill's Game Experiment," 58 *Econometrica* 1065 (1990).

[4] Camerer, Colin, "Individual Decision Making," in Hagel & Roth, *The Handbook of Experimental Economics*, Princeton U. Press (1995).

[5] Camerer, Colin and Teck Hua Ho, "Experience-weighted Attraction Learing in Normal-form Games," 67 *Econometrica* 827 (1999).

[6] Costa-Gomes, Miguel, Vincent P. Crawford and Bruno Broseta, "Cognition and Behavior in Normal-Form Games, An Experimental Study," 69 *Econometrica* 1193 (2001)

[7] Edwards, Ward, "Probability Learning in 1000 Trials," 62 *J. of Experimental Psych.* 385 (1961).

[8] Fudenberg, Drew and David K. Levine, *The Theory of Learning in Games*, MIT (1998).

[9] Gigzerenzer, Todd and the ABC Research Group, *Simple Heuristics that Make Us Smart*, Oxford (1999).

[10] Greene, William H., *Econometric Analysis, 4th Ed.,* Prentice Hall (2000).

[11] Kahneman, Daniel and Amos Tversky, "Subjective Probability: A judgment of representativeness," 3 *Cognitive Psychology* 430 (1972).

[12] McKenzie, Craig, "Judgment and Decision Making," UCSD Working Paper (2002).

[13] Mookherjee, Dilip and Barry Sopher, "Learning Behavior in an Experimental Matching Pennies Game," 7 *Games and Economic Behavior* 62 (1994)

[14] Oechssler, Jorg and Burkhard Schipper, "Can You Guess the Game You Are Playing?," 43 *Games and Economic Behavior* 137 (2003).

[15] O'Neill, Barry, "Nonmetric Test of the Minimax Theoyr of Two-person Zerosum Games," 84 *Proc. Nat. Academy of Sciences,* 2106 (1987)

[16] O'Neill, Barry, "Comments on Brown and Rosenthal's Reexamination," 59 *Econometrica* 503 (1991).

[17] Rabin, Matthew, "Inference by Believers in the Law of Small Numbers," 117 *Quarterly Journal of Economics* 775 (2002).

[18] Rapoport, Amnon and David V. Budescu, "Generation of Random Series in Two-Person Strictly Competitive Games," 121 *J. of Experimental Psychology: General* 352 (1992)

[19] Rapoport, Amnon and David V. Budescu, "Randomization in Individual Choice Behavior," 104 *Psychological Review* 603 (1997).

[20] Roth, Alvin E. and Ido Erev, "Predicting how People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," 88 *Amer. Econ. Rev.* 848 (1998).

[21] Salmon, Timothy C., "An Evaluation of Econometric Models of Adaptive Learning," 69 *Econometrica* 1597 (2001).

[22] Shachat, Jason M., "Mixed Strategy Play and the Minimax Hypothesis," 104 *J. Econ. Theory* 189 (2002).

[23] Shachat, Jason M. and J. Todd Swarthout, "Learning about Learning in Games through Experimental Control of Strategic Interaction," Arizona Working Paper, (2002).

[24] Shachat, Jason M. and J. Todd Swarthout, "Do We Detect and Exploit Mixed Strategy Play by Opponents?" Arizona Working Paper (2003).

[25] Sprent, Peter *Applied Nonparametric Statistical Methods, 2d Ed.*, Chapman & Hall (1995).