



Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb



Edgar Allan Poe's riddle: Framing effects in repeated matching pennies games [☆]

Kfir Eliaz ^a, Ariel Rubinstein ^{b,c,*}^a Department of Economics, Brown University, United States^b The University of Tel Aviv Cafés, Israel^c New York University, United States

ARTICLE INFO

Article history:

Received 16 July 2008

Available online xxxx

JEL classification:

C91

C72

Keywords:

Matching pennies

Even and odd

Framing effects

Edgar Allan Poe

ABSTRACT

Framing effects have a significant influence on the finitely repeated matching pennies game. The combination of being labelled “a guesser,” and having the objective of matching the opponent's action, appears to be advantageous. We find that being a player who aims to match the opponent's action is advantageous irrespective of whether the player moves first or second. We examine alternative explanations for our results and relate them to Edgar Allan Poe's “The Purloined Letter.” We propose a behavioral model which generates the observed asymmetry in the players' performance.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

In his short story “The Purloined Letter,” Edgar Allan Poe writes,

“I knew one about eight years of age, whose success at guessing in the game of ‘even and odd’ attracted universal admiration. This game is simple, and is played with marbles. One player holds in his hand a number of these toys, and demands of another whether that number is even or odd. If the guess is right, the guesser wins one; if wrong, he loses one. The boy to whom I allude won all the marbles of the school. Of course he had some principle of guessing; and this lay in mere observation and admeasurement of the astuteness of his opponents. For example, an arrant simpleton is his opponent, and, holding up his closed hand, asks, ‘are they even or odd?’ Our schoolboy replies, ‘odd,’ and loses; but upon the second trial he wins, for he then says to himself, ‘the simpleton had them even upon the first trial, and his amount of cunning is just sufficient to make him have them odd upon the second; I will therefore guess odd;’ – he guesses odd, and wins. Now, with a simpleton a degree above the first, he would have reasoned thus: ‘This fellow finds that in the first instance I guessed odd, and, in the second, he will propose to himself, upon the first impulse, a simple variation from even to odd, as did the first simpleton; but then a second thought will suggest that this is too simple a

[☆] We thank Marina Agranov for help in running the experiments, Eli Zvuluny for programming, Yoav Benyamini for advising us on statistical matters, Yaniv Ben-Ami and Zhauguo Zhen for their assistance in the data analysis and John Wooders and Jörg Stoye for helpful comments. We wish to especially thank the Associate editor of this journal and three referees for their exceptional comments which helped improve the paper. Financial and technical support from the Center for Experimental Social Science at New York University is gratefully acknowledged.

* Corresponding author at: The University of Tel Aviv Cafés, Israel.

E-mail address: rariel@post.tau.ac.il (A. Rubinstein).

		Player 2	
		L	R
Player 1	T	0,1	1,0
	B	1,0	0,1

Fig. 1.

variation, and finally he will decide upon putting it even as before. I will therefore guess even;’ – he guesses even, and wins. Now this mode of reasoning in the schoolboy, whom his fellows termed ‘lucky,’ – what, in its last analysis, is it? ‘It is merely,’ I said, ‘an identification of the reasoner’s intellect with that of his opponent.’ ”

Poe could have assigned the gifted boy either of the two roles: the first mover, who chooses the number of marbles, or the second mover, who guesses whether that number is odd or even. Poe described the boy as the guesser. Our experimental evidence suggest that this may not have been just a coincidence.

The game Poe describes is essentially a standard (finitely) repeated game in which the stage game is the two-player zero-sum game known as “matching pennies” (see Fig. 1).

All game theoretic solutions we are aware of predict that each player’s chances of “winning” at each round is 50%. Hence, the win rate of each player is distributed according to $Binomial(n, 0.5)$ where n is the number of rounds.

Poe, however, subjects the above game to three framing effects. First, he introduces timing: players move sequentially but without observing each other’s actions. Second, he labels the players: the second mover is referred to as “a guesser.” Third, he labels the actions: one player chooses an odd or even number and another announces “odd” or “even.” A feature of this labeling is that one player wins if his action matches that of his opponent, while the other player wins if the two choose different actions. Do these framing effects have any real impact on the behavior of players? In particular, can such framing effects give an advantage to one of the players, as Poe seems to suggest in his story?

To address these questions, we conducted a series of laboratory experiments in which subjects played different variants of a finitely repeated matching-pennies game. In the baseline treatment (G1), subjects played a simplified version of Poe’s original game, in which the stage game displayed the three framing effects mentioned above: (1) players moved sequentially, but the second player did not observe the action chosen by the first player, (2) the first player is labeled “a misleader” and the second player is labeled “a guesser,” and (3) the actions available to each player are labeled “0” and “1,” such that the guesser wins whenever he chooses the same action as the misleader. Subjects were randomly paired to play 24 rounds of this game (sufficiently many rounds to allow for learning and not excessive to prevent drift of attention).

Our experimental findings in G1 suggest that the combination of being labelled *a guesser*, having the objective of *matching* the opponent’s action and moving *second*, is advantageous. On average, guessers won about 53% of the rounds, and won more points than the misleader in 28 of the 55 participating pairs, while misleaders won more points in only 16 of the pairs (11 pairs ended with a “draw”). Hence, the conjunction of all three framing effects tilts the game in favor of the guesser. The guesser’s advantage is not weakened in the second half of the game, indicating that learning does not eliminate it.

We conjecture that each of the above three framing manipulations could be a source of the guesser advantage in G1.

(1) Past studies in psychology have shown that *priming* of a particular role or personality trait can affect subsequent behavior in the direction of the trait or role (see Bargh et al., 1996). In our context, player 1 is primed to be a misleader and focus more on creating seemingly random patterns and less on predicting his opponent’s response. Player 2, on the other hand, is primed to be a guesser and focus on forecasting his opponent’s action by detecting regularities in the misleaders’ past actions. The guesser’s advantage may be traced to the empirical observation that individuals are not good at randomizing and are better skilled at detecting patterns (see for example, Bar-Hillel and Wagenaar, 1991). This suggests that while attempting to randomize, the misleaders create systematic patterns that are partially detected by the guesser.

(2) The labels on the actions create an asymmetry between the players: one player wins if he matches his opponent’s action, while the other player wins if he mismatches his opponent action. The notion of *stimulus-response compatibility* suggests that it is easier for a subject to match an exogenous stimulus than to mismatch it (see Simon et al., 1981). In our case, the stimulus is a player’s own guess of his opponent’s action. We conjecture that in this case as well, the guesser’s task of matching the stimulus is easier than the misleader’s task of mismatching it.

(3) The role of timing-without-observability is somewhat of a puzzle to us. Previous studies have demonstrated its role in selecting equilibria (see the related literature section). However, it does not seem that those studies are relevant to a zero-sum game with a unique Nash equilibrium.

To better understand the contribution of each framing effect, we conducted four treatments where subjects played variants of G1 that differed only in the framing of the basic stage game. In G2 we dropped the players’ labels, misleader and guesser, and framed the one shot game as an “even and odd” game where player 1 is the “odd” player and player 2 is the “even” player (namely, player 2 wins if the sum of the chosen numbers is even). Our data reveals an advantage to the even player: he won on average 54% of the rounds. This suggests that two of the three framing effects – moving second and having to match the opponent’s action – are enough to generate an advantage.

In order to check whether the advantage of player 2 in G1 and G2 is actually a “second-mover advantage,” we ran two “mirror” treatments: G1*, which is identical to G1 except that it is player 1 who has to guess player 2’s future choice; and G2*, which is identical to G2 except that player 1 is the “even” player. We don’t find a significant advantage to the misleader/odd-player when he moves second.

To further explore the role of timing, we ran a treatment, G3, in which we kept the sequential structure of the original game but dropped the non-neutral labels for *both* the players *and* the actions. In G3, the first player chose a letter “a” or “i,” while the second player chose a letter “s” or “t.” Player 1 (player 2) wins whenever the word which is created is either “at” or “is” (“as” or “it”). In this version, where neither priming nor stimulus-response compatibility played a role, we cannot reject the hypothesis that each player is equally likely to win. This suggests that timing-without-observability alone is not enough to generate an advantage to the second mover.

In light of these findings, we propose a behavioral model, which generates the observed advantage to the player in the role of a guesser/even-player. In this model the players use the following stochastic behavioral rule. With some probability they toss a biased coin to determine whether or not they keep their previous action (this bias can be in favor or against switching). With the complementary probability, they “think strategically” and respond optimally to the prediction of their opponent, which is based on his last action. We allow for two asymmetries in this model. First, we allow for differences in the frequencies by which the two players think strategically. Second, we allow for the possibility of a mistake in executing the best response to the prediction. We assume that: (i) the guesser/even player is more likely to think strategically than the misleader/odd-player, and (ii) the misleader/odd-player is more likely to err in executing his optimal response since it is a more complicated operation (he has to mismatch where the guesser/even player has to match). We find that both asymmetries, each by itself and to a larger extent together, generate about 52% advantage to the guesser/even-player for a reasonable range of parameters.

Related literature. Closest in spirit to our paper is Rubinstein et al. (1996), which examines several variations of a hide and seek game. In the baseline version, “the hider” has to “put a treasure” in one of four boxes labeled from left to right A–B–A–A, and “the seeker” has the opportunity to search for it in one of the boxes. According to the game-theoretical prediction of this zero-sum game, both players uniformly randomize between the four boxes and thus, the probability that the seeker finds the treasure is 25%. In the original research, as well as in subsequent replications of the experiment, central A is the most popular choice among both the hiders and the seekers. For example, data collected through the didactic website, <http://gametheory.tau.ac.il>, reveals that the distribution of 1,129 misleaders was (17%, 24%, 36%, 22%) and the distribution of 1,116 guessers was (10%, 26%, 47%, 16%). According to these distributions, the probability of finding the treasure is 28.6%, significantly above the expected 25%. It seems that subjects in both roles, who wished to play randomly, looked for a choice which is not prominent. They excluded the “distinctive” B and the two boxes at the edges, leaving them with the unique “nondistinctive” box, central A.

However, it is not clear whether this result may be interpreted as evidence for the existence of a guesser’s advantage since this data cannot distinguish between the following alternative explanations: (i) the normal choice of A is close to 36% and many guessers think strategically, or (ii) the normal choice rate of A is closer to 47% and many of the misleaders think strategically. In contrast, we *do* claim to demonstrate a guesser’s advantage since the data do not point at a prominent action in our setting.

Also related is Attali and Bar-Hillel (2003), which provides field evidence on the difficulty in preventing guessers from exploiting regularities in supposedly random sequences of play. The authors looked at high-stakes tests such as SATs and discovered that the creators of these tests had a policy of balancing, rather than randomizing, the answer keys. This policy was vulnerable to significant exploitation by examinees. Note, however, that the exam situation is not truly a zero-sum game: the composers of the exam want to prevent examinees from guessing too successfully, but they do not want to mislead and fail examinees too often in cases where they do not know the answer.

Repeated matching pennies games were examined in the literature by Mookherjee and Sopher (1994) and Ochs (1995). In those experiments the matching pennies game was presented to the players as a bi-matrix. No significant deviation from the mixed strategy equilibrium was reported. Shachat and Wooders (2001) shows that in games like repeated matching pennies, the Nash equilibrium is the repeated play of the stage game equilibrium, even if the players are not risk neutral.

O’Neill (1987) provides a test of the maxmin prediction in a zero-sum game. In his game each player has to choose one of the four actions, *a*, *b*, *c* or *d*, and player 1 wins if the two players choose distinct actions other than *d*, or if both choose *d*. Both players’ maxmin mixed strategies are (0.2, 0.2, 0.2, 0.4), and the equilibrium winning probability of player 1 is 40%. O’Neill claimed that in the finitely repeated game the equilibrium prediction is confirmed. The claim was challenged by Brown and Rosenthal (1990) who observed that the data at the individual level reveals significant serial correlations in each player’s choices. According to data collected by one of us via the site <http://gametheory.tau.ac.il>, the one shot game behavior is significantly different than the maxmin strategies, (0.16, 0.18, 0.2, 0.45) and (0.14, 0.18, 0.25, 0.44) for player 1 and player 2, respectively, though the probability of winning by player 1 is 0.403, amazingly close to the equilibrium prediction. In contrast, framing has no effect in the one shot matching pennies.

Several previous experimental studies have demonstrated that timing matters even when two players move sequentially and the early move is unobservable by the second player. All of these papers have focused exclusively on the role of timing in selecting among multiple Nash equilibria. Cooper et al. (1993) have shown that in coordination games players tend to coordinate on the equilibrium preferred by the first mover. Weber et al. (2004) test the notion of “virtual observability,”

Table 1

Key features of the five treatments.

Treatment	Player 1	Player 2	1's actions	2's actions	Player 1 wins if:
G1	Misleader	Guesser	0/1	0/1	Player 2 does not guess him correctly
G1*	Guesser	Misleader	0/1	0/1	He predicts correctly player 2's choice
G2	Odd	Even	0/1	0/1	The sum of chosen numbers is odd
G2*	Even	Odd	0/1	0/1	The sum of chosen numbers is even
G3	–	–	a/i	s/t	The word created is "at" or "is"

Table 2

Summary of the results.

		<i>n</i> = 55	<i>n</i> = 29	<i>n</i> = 29	<i>n</i> = 26	<i>n</i> = 27
Player 2's role		Guesser	Misleader	Even	Odd	As/it
Player 2's wins	<i>B</i>(24, 0.5)	G1	G1*	G2	G2*	G3
0–7	4%	0%	3%	0%	0%	0%
8	4%	0%	0%	3%	8%	11%
9	8%	7%	7%	3%	8%	7%
10	12%	7%	7%	3%	19%	4%
11	15%	15%	14%	21%	27%	11%
12	16%	20%	21%	14%	8%	22%
13	15%	16%	14%	17%	12%	11%
14	12%	11%	21%	10%	8%	22%
15	8%	20%	10%	10%	8%	4%
16	4%	4%	3%	7%	0%	0%
17–24	4%	0%	0%	10%	4%	7%
p-value (1-sample K-S*)		1.51%	4.14%	7.78%	46.69%	7.26%
Player 2's win rate	50.0%	52.6%	51.4%	54.2%	47.9%	51.1%
Player 2's mean score (standard deviation)	12.00	12.61 (0.26)	12.34 (0.4)	13.00 (0.45)	11.50 (0.44)	12.26 (0.51)
p-value (normal test)		3.1%	22.4%	1.4%	14.9%	29.1%

* Bootstrapped to correct for discreteness of the distributions.

according to which players behave as if moves were observable, and hence, tend to coordinate on a Nash equilibrium, which is also the subgame-perfect equilibrium of the sequential game with observable actions. Our game has a unique Nash equilibrium and thus, the effect of timing in our experiment is completely different than its effect in those previous studies.

Finally, let us also mention the game theoretical models which analyze the repeated matching pennies game using the bounded rationality approach. (See for example Ben-Porath, 1993, Neyman, 1997, and Guth et al., 2005.)

2. Experimental design

Subjects were recruited from a pool of students at New York University. The experiment was conducted at the laboratory of the Center for Experimental Social Science (C.E.S.S.) at NYU. An even number (at least 8, at most 20) of subjects participated in each session. The computer randomly matched the subjects into anonymous pairs (i.e., we used a "partners" design). Each pair played a 24-period repeated matching pennies game; in each round, player 1 moved first and player 2 moved second *without observing* player 1's choice. During each round, each player could see on his screen the choices made by the players in previous rounds and the number of points accumulated so far. We examined five versions of the game that differ only in the framing of the game. Appendix A contains the detailed instructions for the basic treatment, G1. The instructions for the other treatments are provided in Section 3. In each session all subjects played the same version of the game. At the end of each session the experimenter calculated the rewards: a sum of \$5 was paid to a subject for participation and additional 50 cents were given per each win in the stage game. Thus, the average amount paid to subjects is \$11. The entire process took about 30 minutes.

Table 1 outlines the key features of each of the treatments. A full description of each treatment appears in Section 3. In each of the first four games, a bold face letter indicates the player who wins if the two actions coincide.

3. The basic findings

Table 2 summarizes the results of the five treatments. The results of each treatment are presented in a separate column and include the distribution of points gained by the second movers, their win rate, and both the mean and standard deviation of their score. The results are compared with the distribution of $B(24, 0.5)$. The first statistical test is the one-sample Kolmogorov–Smirnov, which is used to compare a sample with a reference probability distribution (in our case $B(24, 0.5)$). The last line presents the p-values generated by the normal test. The null hypothesis in G1, G2 and G3 is that the mean score of player 2 is at most 12 whereas in G1* and G2* the hypothesis is that the mean score is at least 12.

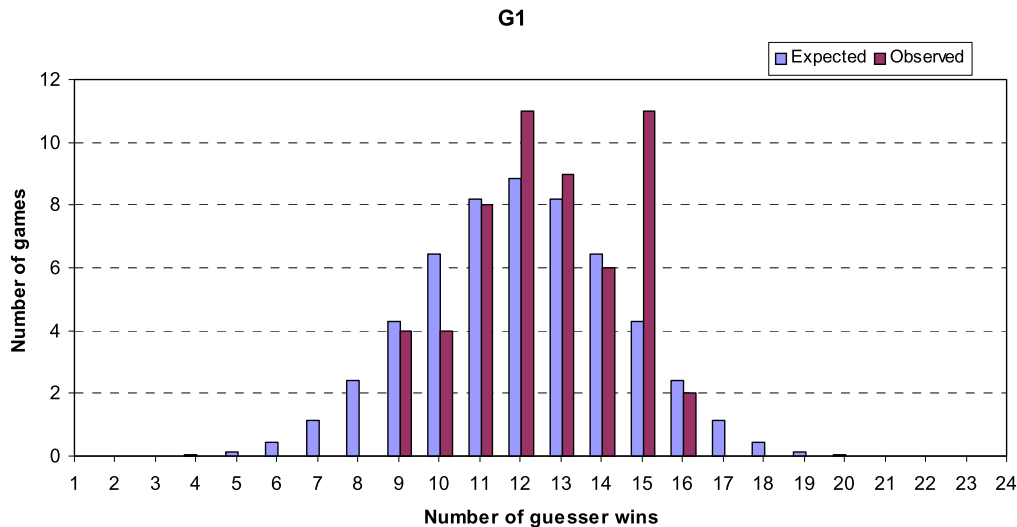


Fig. 2. Actual vs. predicted distribution of scores in G1.

3.1. The misleader-guesser game

The baseline treatment is:

G1: Imagine that you are about to play a 24-round game. In the game there are two players, a misleader and a guesser. At each round, the misleader moves first and has to choose a number, 0 or 1. The guesser moves second and has to guess the misleader's choice. The guesser gets a point if his guess is correct and the misleader gets a point if the guesser fails. Each player's aim is to get as many points as possible. At each round each player can see on the screen the entire history of the game up to that point.

In G1 the guesser has a significant advantage. On average, guessers gained 12.6 points, namely, they won 53% of the rounds. In the repeated game, the guesser won more points than the misleader in 28 of the 55 games, while the misleader won more points in only 16 games (11 repeated games ended with a "draw"). The results are significant at the 3.1% level.

To get a better sense of the guesser's advantage it is helpful to compare the distribution of our subjects' points in G1 with the distribution predicted by $B(24, 0.5)$. This is depicted in Fig. 2.

One may think that a success rate of 0.526 is an insignificant deviation from the 0.5 prediction. We disagree. First note that even if both players would use independent mixed strategies, which assign probability 0.6 to one of the two actions, the expected guesser's win rate would be only 0.52 ($0.6 \times 0.6 + 0.4 \times 0.4$). This indicates that subjects' behavior is far from the game theoretical prediction, and therefore, sophisticated players may obtain an even higher win rate. Second, this "slight" advantage of the guesser could become inflated for a variation of the game (resembling the "even and odd" game in Edgar Allan Poe's "The Purloined Letter," which is discussed later) where the winner is the player who gets first to an advantage of some fixed number of points.

As mentioned previously, G1 introduces three major framing effects to a standard repeated matching-pennies game: being labelled a guesser/misleader, aiming to match/mismatch the opponent's action and moving second/first. To investigate the extent to which these framing effects may have contributed towards the guesser's advantage, we tested several variations of G1. We discuss these in the next subsections.

3.2. The even and odd game

To control for non-neutral labeling of players' roles, we ran the following treatment:

G2: You are about to play a 24-round game. There are two players in the game: player 1 and player 2. At each round, player 1 moves first and player 2 moves second. Each player has to choose a number, 0 or 1. Player 1 gets a point if the sum of the chosen numbers is odd. Player 2 gets a point if the sum is even. Each player's aim is to get as many points as possible. At each round each player can see on the screen the entire history of the game up to that point.

G2 is a modification of G1 in which the players are not labeled as a misleader and a guesser. While player 2 has to simply choose his guess of player 1's action, player 1 has to choose the *opposite* of what he anticipates player 2 will choose.

Our experimental results hint at an advantage of the even player (player 2) who won 54% of the games (on average 13 points), a significant deviation from the game theoretical prediction with $p = 1.4\%$. This suggests that two of the framing effects – aiming to match the opponent's action and moving second – are enough to generate a significant guesser advantage.

3.3. The order effect

The results of G1 and G2 leave the possibility that what we observe is actually a “second-mover advantage.” In order to examine this possibility, we investigated two other treatments, dual to G1 and G2 in the sense that the roles of the two players are exchanged. G1* is a guesser–misleader game in which player 1 guesses player 2's choice. G2* is an even and odd game in which player 1 is the even player.

In both dual game, the second mover (the misleader in G1* and the odd-player in G2*) obtains no advantage. In the even and odd game (G2*), the even player who is there player 1, even obtained an advantage and gained on average 12.5 points, however this advantage is significant with $p = 15\%$ only.

3.4. The “as/it” game

The results of G1* and G2* still do not address the question of whether timing by itself (i.e., keeping the labeling of actions and roles as neutral as possible) can generate a guesser's advantage. We address this question in the final treatment (G3) where the order of moves is the only asymmetry between the two players.

G3: You are about to play a 24-round game. In the game there are two players: player 1 and player 2. At each round, the two players compose a two-letter word. Player 1 moves first and has to choose the first letter of the word, a or i. Player 2 moves second and has to choose the second letter of the word, s or t. Player 1 gets a point if the word formed is either at or is. Player 2 gets a point if the word formed is either as or it. Each player's aim is to get as many points as possible. At each round each player can see on the screen the entire history of the game up to that point.

The distribution of points among the 27 pairs who played the game was close to that of the *Binomial*(24, 0.5). In particular, player 2 won only 51.1% of the rounds (331 out of 648 rounds), an insignificant advantage. We interpret these findings as evidence that timing-without-observability on its own cannot generate an advantage to one of the players.

To summarize, our results suggest that an appropriate choice of framing can have an asymmetric effect in what may be perceived as a totally symmetric game. Our data indicates that being a player who aims to match the opponent's action is advantageous irrespective of whether the player moves first or second.

4. Searching for explanations of the guesser's advantage

4.1. Is there a prominent action?

The following could be a possible explanation for the guesser's advantage in G1. Assume there was a prominent action, which is more likely to be chosen by subjects in both roles. Then given the payoffs of the game, the tendency of subjects to choose the prominent action would give an advantage to the guesser (this explanation is similar to the one given in Rubinstein et al. (1996) for the seeker's advantage).

However, according to the data no action is prominent for *both* players. In the first round of G1, 38% of the misleaders and 67% of the guessers chose “1.” Taking these numbers as estimates for the players' strategies yields a prediction that the guesser wins at each round with probability 0.46, which does not explain the guesser's advantage (similar results were obtained for the other treatments).

Furthermore, a prominent action does not show up in the data of the overall play of the game (chi-square test yields a p-value of 0.8). The following table presents the overall frequencies of the outcomes in all rounds in G1. The misleader chooses slightly more zeros and the guesser chooses slightly more ones. Using these frequencies to estimate mixed strategies for the whole game yields a prediction that the guesser wins with probability 0.4996 ($= 0.492 \times 0.523 + 0.508 \times 0.477$), significantly below the 0.525 frequency of the guesser's wins in G1 (see Table 3).

Thus, the existence of some prominent action does not explain the guesser's advantage.

4.2. Is the guesser's advantage an artifact of inexperienced players?

One might conjecture that the deviations from the game theoretic prediction are limited to the beginning of the game and that with time, players learn to play the mixed strategy equilibrium. By this approach, one would expect that the mean scores in the second half of the game would be more symmetric than the mean scores in the first half. The data summarized in Table 4 does not confirm this line of thinking. The last column of the table presents the results of a two-sample K-S tests of the hypothesis that the distributions of the results in the two halves is drawn from the same distribution. The test clearly supports this hypothesis.

Table 3
Distribution of one shot game outcomes.

		Guesser:		
		0	1	
Misleader:	0	27.0	25.2	52.3
	1	22.2	25.5	47.7
		49.2	50.8	

Table 4
Mean number of rounds won by the guesser.

Game	Guesser's mean score			Two-sample K-S* (p-value)
	In the full game	In the first half	In the second half	
G1	12.6	6.2	6.4	0.966
G1*	11.7	5.9	5.8	0.875
G2	13.0	6.5	6.5	0.859
G2*	12.5	6.3	6.2	0.987
G3	12.3	6.2	6.0	0.985

* Bootstrapped to correct for discreteness of the distributions.

Table 5
Transition probabilities of one shot outcomes.

$n = 1,265$	(1, 1)	(0, 0)	(1, 0)	(0, 1)
(1, 1)	35.1%	17.7%	21.1%	26.1%
(0, 0)	16.9%	36.1%	24.1%	23.0%
(1, 0)	25.0%	28.2%	24.6%	22.1%
(0, 1)	25.4%	26.7%	21.0%	27.0%
Total	25.5%	27.3%	22.7%	24.6%

Table 6
Persistence rate after success/failure.

	Repeats action
Misleader success	51.8%
Misleader failure	57.7%
Guesser failure	52.6%
Guesser success	60.7%

4.3. Persistence effect

Applying the Wald–Wolfowitz runs test to the individual data does not reject the hypothesis that a vast majority of the subjects do produce a random sequence. In G1, for example, the hypothesis of a random sequence is not rejected ($p = 5\%$) regarding 87% of players 1 and 76% of players 2.

One might think that we would observe a “negative recency effect”, as reported in the psychological literature regarding the production of random sequences (see for example, Kahneman and Tversky, 1972). According to this literature, when producing a random sequence of zeros and ones, people tend to change their choices from period to period more than 50% of the time. Our data shows the opposite. We observe a tendency for persistence. Players 1 and 2 in G1 repeat their past choices in 55% and 57% of the cases, respectively. Similar frequencies (between 53% and 60%) are observed in the other versions of the game.

Note that the persistence effect does not explain the guesser's advantage since it could be that the players are equally likely to repeat their choices after a misleader's win as after a guesser's win.

4.4. A partial explanation: asymmetry in response to past success

There is a transition matrix of the one shot outcomes in G1 in Table 5. The one striking observation is the exceptionally significant tendency for persistence of the outcomes (1, 1) and (0, 0) and the low transition probabilities from (1, 1) to (0, 0) and from (0, 0) to (1, 1). To shed more light on the transition matrix, let us look at Table 6, which displays the chances that

each of the players repeats his action following a success and following a failure. This table reveals that the guesser tends to repeat his action after success and the misleader tends to retain his past action after failure. Suppose the guesser followed a simple strategy where he repeats his past action with probability 0.61 after success and probability 0.53 after failure, and the misleader repeats his action with probability 0.58 after failure and with probability 0.52 after success. Then a guesser's success will be followed by another success with probability $0.58 \times 0.61 + 0.42 \times 0.39 = 0.52$, and a misleader's success will be followed by an equally likely win by each of the players. Calculating the stationary distribution of the Markovian matrix yields a 51% success rate for the guesser, and thus, provides a partial explanation for the guesser's success.

The above is the best partial explanation we found in the data for the guesser's advantage. It should be noted that this explanation *does not* extend to G2. There, player 1 (the odd player) repeats his action after success with probability 0.63 and after failure with probability 0.50, whereas player 2 (the even player) repeats his action after failure with probability 0.56 and after success with probability 0.53. These numbers would be more consistent with a slight advantage for player 1, while in G2 we observe an advantage for player 2.

4.5. One more curious fact

In G1, there were 73 histories in which the same outcome, (1, 1) or (0, 0), was repeated in the last $k \geq 3$ periods. In 55% of those cases the guesser succeeded also in the next period. Strikingly, there were only 28 histories in which the same outcome, (1, 0) or (0, 1), was repeated in the last $k \geq 3$ periods. Such histories were followed by yet another victory for the misleader in only 36% of the cases. This asymmetry between the players may reinforce our intuition that guessers were more alert than misleaders.

5. Modeling the players' behavior

There are two potential asymmetries that may explain the guesser's/even-player's advantage.

- (i) The misleader/odd-player has to execute a more complicated mental operation. The guesser/even-player has to predict his opponent's move and to choose the same action. In contrast, the misleader/odd-player has to predict his opponent's move and to choose the *opposite*, an operation which is more likely to yield a mistake.
- (ii) The guesser/even-player is more primed than the misleader/odd-player to think strategically, namely to infer the next move from the last move and to respond optimally to his prediction.

The following model attempts to capture the above asymmetries. Consider a repeated matching-pennies game, where it is player 2 who wins a stage game if the choices match. Each player chooses his initial action at random. From then on, player i 's strategy is given as follows. With probability $1 - \beta_i$ he randomizes and with probability β_i he "thinks strategically." When a player randomizes he has a bias. His last action can predict his next action with probability $\alpha > \frac{1}{2}$. Here we assume that a player *repeats* his previous action with probability α and switches an action with probability $1 - \alpha$. If we were to assume that a player is more likely to *switch* an action when he tries to randomize (what seems consistent with other experiments), our results will not change, provided that we also adjust the second (strategic) part of player's behavioral

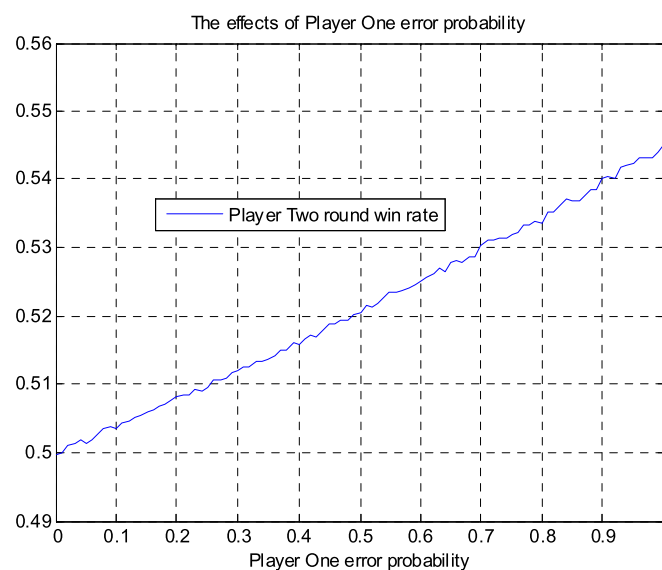


Fig. 3. Player 2's win rate as a function of ε for $\beta_1 = \beta_2 = 0.3$.




Fig. 4. Player 2's win rate in the 24 round repeated game as a function of ϵ for $\beta_1 = \beta_2 = 0.3$.

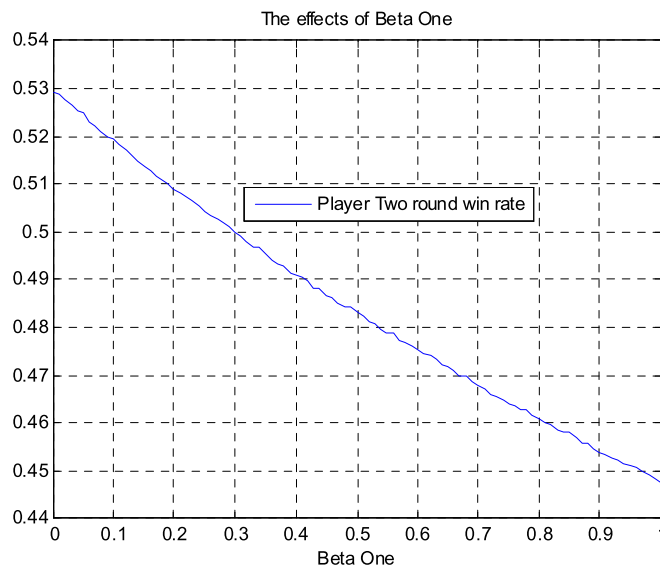


Fig. 5. Player 2's win rate as a function of β_1 with $\beta_2 = 0.3$ and $\epsilon = 0$.

rule. When a player “thinks strategically” he believes that it is more likely that his opponent will repeat his last action and he best responds to this prediction. However, while player 2 executes the best response accurately, player 1 makes a mistake with probability ϵ , and instead of mismatching his opponent's action he matches it.

Assume first that the matching-pennies game is played for infinitely many periods. There are two states: either the players match their actions, or they mismatch. Computing the Markov stationary probabilities for these two states, yields the following expression for player 2's win rate:

$$\frac{[(1-\beta_1)\alpha + \beta_1(1-\epsilon)][\beta_2 + (1-\beta_2)(1-\alpha)] + (1-\beta_2)\alpha[\beta_1\epsilon + (1-\beta_1)(1-\alpha)]}{1 + [(1-\beta_1)\alpha + \beta_1(1-\epsilon)][\beta_2 + (1-\beta_2)(1-\alpha)] + (1-\beta_2)\alpha[\beta_1\epsilon + (1-\beta_1)(1-\alpha)] - [(1-\beta_1)\alpha + \beta_1\epsilon][\beta_2 + (1-\beta_2)\alpha] - (1-\beta_2)(1-\alpha)[(1-\beta_1)(1-\alpha) + \beta_1(1-\epsilon)]}$$

To better understand the relation between this expression and our data, assume next that the matching-pennies game was repeated for only 24 periods. Assume also that $\alpha = 0.6$. We simulated the above model to test whether it can generate the guesser/even-player's advantage observed in our data. Fig. 3 plots player 2's win rate as a function of ϵ for $\beta_1 = \beta_2 = 0.3$ (the grid on this graph is of 1% and each point averages 100,000 runs of the repeated game).

As evident from the figure, the asymmetry in the error's rate (no error for player 2 and positive error rate for player 1) is capable of explaining the guesser/even-player's advantage: when player 1's error lies in the interval [0.2, 0.3], player 2

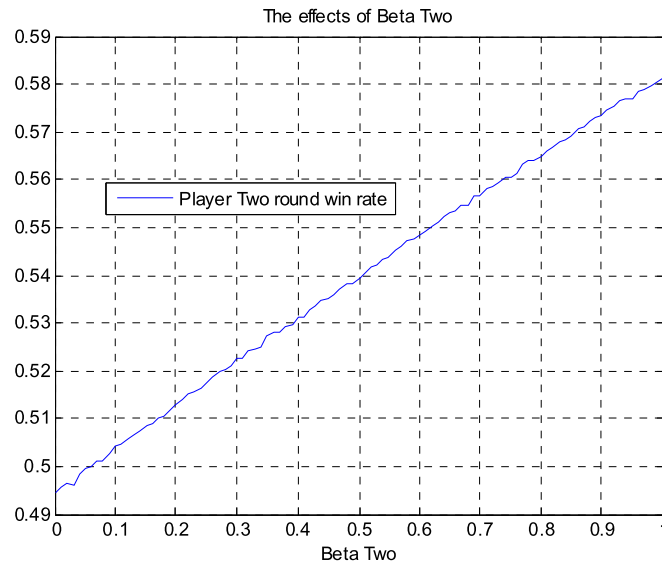


Fig. 6. Player 2's win rate as a function of β_2 for $\varepsilon = 0.2$ and $\beta_1 = 0.1$.




Fig. 7. Player 2's win rate as a function of ε in the alternative model.

gets a lead of 1% to 2% above player 1's win rate. Fig. 4 demonstrates that the effect of this asymmetry is amplified when we consider a player a winner in the repeated game if he scores more than 12 points. Using the same parameters we get that the difference between player 2's win rate and player 1's win rate increases quickly with ε and reaches 10% when the mistake probability is 0.3.

Looking for the effect of the priming asymmetry, we simulated the model without mistakes ($\varepsilon = 0$) and with asymmetry in the likelihood of behaving strategically. Fig. 5 plots player 2's win rate as a function of β_1 when $\beta_2 = 0.3$. As evident from the graph, player 2 has an advantage as long as β_1 is below β_2 . In particular when $\beta_1 = 0.1$ player 2's win rate is 52% (which is the rate we obtain experimentally in G1).

Combining the two asymmetries yields a greater advantage for player 2. Fig. 6 plots the win rate of player 2 as a function of β_2 when $\varepsilon = 0.2$ and $\beta_1 = 0.1$. As we can see, player 2 obtains a win rate of 0.52 when β_2 is around 0.27.

Note that some other variants of our model provide paradoxical results. Consider, for example, another behavioral model for the 24-period repeated matching-pennies game. As before, player i 's behavioral rule has two components, random and strategic. When a player randomizes he "uses" a random device, which has a systematic bias either towards the action "1" or towards the action "0"; the device yields the outcome 1 with either probability 0.6 or 0.4. The bias is determined at the beginning of the game and is fixed throughout the 24 rounds. In contrast to the first model, the strategic component

is triggered by the behavior of his opponent. When a player notices that the other player played the same action twice in a row, he best responds to that action. The only asymmetry we incorporate to this model is that player 1, and only player 1, makes a mistake with probability ε when attempting to best respond to an action that player 2 repeated twice.

Note that for this alternative model it is much more complicated to derive the Markov stationary probabilities (assuming infinitely many repetitions of the stage game) as there are many possible states. We, therefore simulate player 2's win rate as function of player 1's error rate, assuming for 24 repetitions of the stage game. The results are plotted in Fig. 7. Surprisingly, player 1's mistake works in his favor (as long as $\varepsilon < 0.6$)! This fact has to do with the structure of the game, which creates artificial synchronization between the timing in which the players best respond to each other. This synchronization does not arise in the first behavioral model where it is rare that the two players will think strategically in the same period.

6. Concluding remarks

6.1. Back to Edgar Allan Poe

Poe gives the cunning boy the role of the guesser. The boy distinguishes between an opponent who is simple minded and one who is not. If the opponent is simple minded (in our modern jargon, he uses level zero reasoning), the boy predicts that he will not change his action after success and thus, by changing his own action, he wins the next round. If the opponent is not simple minded (i.e., he uses level one reasoning), the boy predicts that he will change his action after success, expecting the boy to do the same. The cunning boy is even more sophisticated (i.e., he uses level two reasoning) and therefore beats the level one opponent.

Our findings are related to Poe's intuition and would fit nicely in our explanation of subjects' behavior after a guesser's success (rather than a guesser's failure), assuming that guessers believe that the misleaders are not simple minded. As we described in Section 4.4, following a misleader's failure, 57.5% of the misleaders and 60.7% of the guessers repeat their action. However, Poe demonstrates the boy's cunningness after his failure where we do not find traces of such thinking.

Note that the game Poe describes is somewhat different from ours. The closest parallel will be a version in which the first player to have an advantage of k points wins the match. If the guesser's victories are i.i.d. across games with a winning probability of q , then his chances of being the first to have an advantage of k points is

$$\frac{1}{1 + \left(\frac{1-q}{q}\right)^k}.$$

It follows from this formula that guessers' "minor" advantage in the one shot game is translated to a huge advantage in Poe's game. For example, if $q = 0.526$ and $k = 4$ the guesser will win the game with probability 0.6, or if $k = 8$ his chances of winning will jump to 0.7.

6.2. Game Theory

Much has been written about the failures of Game Theory to predict behavior. Advocates of Game Theory question the significance of the experimental evidence which reject the game theoretical predictions. Some claim that the refutation was conducted over games which the subjects were not experienced to play. Some argue that Nash equilibrium fits only situations where the game is played often enough so that a convention is settled. Our experiment should have been the ideal stage for Game Theory to "work." A simple and familiar game, very simple equilibrium strategies, and repeated long enough to allow players to converge to equilibrium behavior. Nevertheless, our findings hint at a systematic deviation from the Game Theoretical prediction. More importantly, it hints at the possibility that even in a simple situation like matching pennies, the framing of the problem primes players to use different procedures for playing the game, and that these procedures do not perform equally well.

Appendix A. Instructions

Introduction. This is an experiment in interactive decision-making. Various research institutes have provided the money for this experiment and you can make a considerable amount of money in a short time if you pay attention.

The experiment. You are about to play a 24-round game. In the game, two players: a **misleader** and a **guesser**. At each round, the **misleader** moves *first* and has to choose a number, **0** or **1**. The **guesser** moves *second* and has to guess the **misleader's** choice.

The guesser gets a point if his guess is correct and the misleader gets a point if the guesser fails. Each player's aim is to get as many points as possible. At each round each player can see on the screen the entire history of the game up to that point.

The computer will randomly match the participants in this experiment into pairs. In each pair, one participant will be randomly assigned to play the role of the **misleader**, while the other participant will be assigned the role of the **guesser**.

Payoffs. Each point that you earn in the game will be converted to \$0.50. Your total payoff in the experiment will be equal to the total sum of points you earn in the 24 periods of the game, multiplied by \$0.50, plus the show-up fee of \$3. The minimum payoff is \$3 and the maximum is \$15.

How to start the experiment?

On the screen before you type in the following information:

Course Number: 1323.

E-mail: Look at the number of the station, you are sitting in. If you are sitting in lab 1, the e-mail you should type in is **1lab1@nyu.edu**. Similarly, if you are sitting in lab 2, your e-mail is **1lab2@nyu.edu**. In general, if you are sitting in lab x, your e-mail is **1labx@nyu.edu**.

Password: If you are sitting in lab x, your password is **1x**. For example, if you are sitting in lab 1, your password is 11. If you are sitting in lab 13, your password is 113, and so forth.

References

- Attali, Yigal, Bar-Hillel, Maya, 2003. Guess where: The position of correct answers in multiple-choice test items as a psychometric variable. *J. Educat. Measurement* 40, 109–128.
- Bargh, John A., Chen, Mark, Burrows, Lara, 1996. Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *J. Person. Soc. Psych.* 71, 230–244.
- Bar-Hillel, Maya, Wagenaar, Willem A., 1991. The perception of randomness. *Adv. Appl. Math.* 12, 428–454.
- Ben-Porath, Elchanan, 1993. Repeated games with finite automata. *J. Econ. Theory* 59, 17–32.
- Brown, James N., Rosenthal, Robert W., 1990. Testing the minimax hypothesis: A re-examination of O'Neill's game experiment. *Econometrica* 58, 1065–1081.
- Cooper, Russell, DeJong, Douglas V., Forstythe, Robert, Ross, Thomas W., 1993. Forward induction in the Battle-of-the-Sexes games. *Amer. Econ. Rev.* 83, 1303–1316.
- Guth, Werner, Kareev, Yaakov, Kliemt, Hartmut, 2005. How to play randomly without random generator – the case of maximin players. *Homo Oeconom.* 22, 231–255.
- Kahneman, Daniel, Tversky, Amos, 1972. Subjective probability: A judgment of representativeness. *Cognitive Psych.* 3, 430–454.
- Mookherjee, Dilip, Sopher, Barry, 1994. Learning behavior in an experimental matching pennies game. *Games Econ. Behav.* 7, 62–91.
- Neyman, Abraham, 1997. Cooperation, repetition and automata. In: Hart, S., Mas-Colell, A. (Eds.), *Cooperation: Game-Theoretic Approaches*. In: NATO ASI Ser. F, vol. 155. Springer-Verlag, pp. 233–255.
- Ochs, Jack, 1995. An experimental study of games with unique mixed strategy equilibria. *Games Econ. Behav.* 10, 202–217.
- O'Neill, Barry, 1987. Non-parametric test of the minimax theory of two-person zerosum games. *Proc. Natl. Acad. Sci. USA* 84, 2106–2109.
- Rubinstein, Ariel, Tversky, Amos, Heller, Dana, 1996. Naive strategies in zero-sum games. In: Guth, W., et al. (Eds.), *Understanding Strategic Interaction – Essays in Honor of Reinhard Selten*. Springer-Verlag, pp. 394–402.
- Shachat, Jason, Wooders, John, 2001. On the irrelevance of risk attitudes in repeated two-outcome games. *Games Econ. Behav.* 34, 342–363.
- Simon, Richard, Sly, Paul E., Vilapakkam, Sivakumar, 1981. Effect of compatibility of S–R mapping on reactions towards the stimulus source. *Acta Psycholog.* 47, 63–81.
- Weber, Roberto A., Camerer, Colin F., Knez, Marc, 2004. Timing and virtual observability in ultimatum bargaining and 'weak link' coordination games. *Exper. Econ.* 7, 25–48.