

# Minimax across a population of games

Ido Erev<sup>1</sup> · Alvin E. Roth<sup>2</sup> · Robert Slonim<sup>3</sup>

Received: 18 June 2015 / Revised: 3 October 2016 / Accepted: 5 October 2016 /  
Published online: 22 October 2016  
© Economic Science Association 2016

**Abstract** Most economic experiments designed to test theories carefully choose specific games. This paper reports on an experimental design to evaluate how well the minimax hypothesis describes behavior across a *population* of games. Past studies suggest that the hypothesis is more accurate the closer the equilibrium is to equal probability play of all actions, but many differences between the designs makes direct comparison impossible. We examine the minimax hypothesis by randomly sampling constant sum games with two players and two actions with a unique equilibrium in mixed strategies. Only varying the games, we find behavior is more consistent with minimax play the closer the mixed strategy equilibrium is to equal probability play of each action. The results are robust over all iterations as well as early and final play. Experimental designs in which the game is a variable allow some conclusions to be drawn that cannot be drawn from more conventional experimental designs.

**Keywords** Game theory · Experimental design · Equilibrium

**JEL Classification** C9

---

**Electronic supplementary material** The online version of this article (doi:[10.1007/s40881-016-0029-3](https://doi.org/10.1007/s40881-016-0029-3)) contains supplementary material, which is available to authorized users.

---

✉ Robert Slonim  
robert.slonim@sydney.edu.au

<sup>1</sup> The Technion, Haifa, Israel

<sup>2</sup> Stanford University, Palo Alto, California, USA

<sup>3</sup> School of Economics, The University of Sydney, Sydney 2006, Australia

## 1 Introduction

Economic experiments usually evaluate the predictions of a theory on a game or task carefully chosen to be a good test of some aspect of the theory. This strategy permits researchers to conclude that either the theory correctly predicts the particular behavior tested or that the theory is not true in general. If the theory is false, the experimental results often suggest how the theory can be modified and/or how the theory needs to be qualified. However, this strategy does not provide information on whether the qualifications (i.e., the reasons a theory is rejected) are robust to similar, yet not identical conditions. For instance, if the results of a low-stakes experiment using specific parameters of a game reject a theory using a student subject population, the theory can be rejected as a general theory, yet the results do not indicate whether the theory would describe behavior among other populations such as professional traders, under other conditions such as higher stakes, or with different parameters. Researchers have therefore studied the robustness of theories across stakes and subject populations (e.g., Slonim and Roth 1998; Roth et al. 1991). But there has been little research focused on the variation of results over a *sample population of games*.<sup>1</sup>

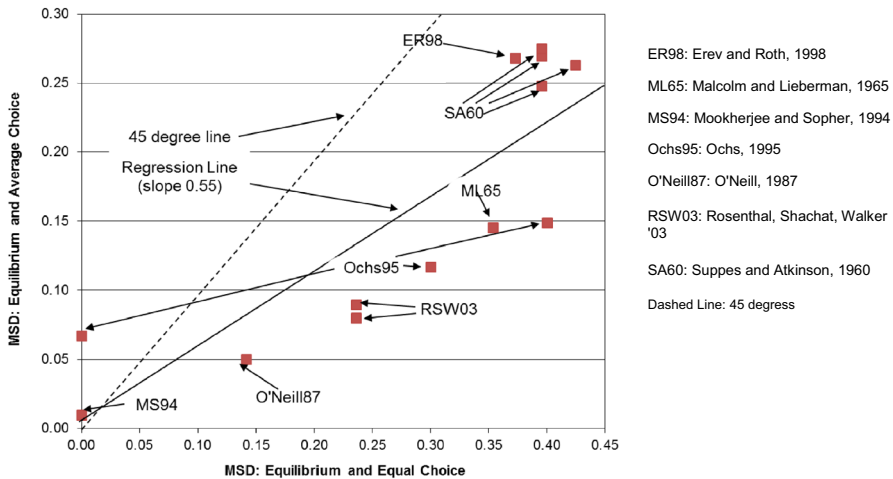
This paper uses a sample of games to examine the accuracy of the minimax hypothesis. A sample of games permits a test of the robustness of the minimax hypothesis—does it fail rarely or often? More importantly, the sample of games permits testing correlations between the magnitude of deviations from theory and the different parameters of the games. These correlations may help explain why past experiments prompt different conclusions regarding the accuracy of a theory, as is the case with the minimax hypothesis.

Past experiments testing the precise predictions of the minimax hypothesis find mixed evidence. Some experiments find that the minimax hypothesis has little descriptive accuracy (Suppes and Atkinson 1960; Malcolm and Lieberman 1965; Erev and Roth 1998), while others find more accuracy (O'Neill 1987; Mookherjee and Sopher 1994; Ochs 1995). While the minimax hypothesis makes many specific predictions (regarding, e.g., serial correlation, joint choice distributions), we focus here on deviations in the observed aggregated proportion of play of each strategy across the population and rounds of play.

Figure 1 shows the mean square distance between minimax play and subject behavior for the past studies as a function of the mean square distance between minimax play and the equal choice mixture. The figure shows that as minimax play gets further from the equal choice mixture across games, minimax play also gets further from subject behavior. OLS regressions show a significant positive relationship between (1) the mean square distance between minimax play and subject choices and (2) the mean square distance between minimax play and equal choice; OLS regressions indicate a slope of 0.55 (standard error 0.10) and we can reject at the  $p < 0.01$  level that the slope equals 0 or 1.<sup>2</sup> However, these results were

<sup>1</sup> Most exceptions (e.g., Choi et al. 2007; Erev et al. 2002, 2007, 2010a, b) use random design to facilitate model comparison. The current research uses part of the data collected in these studies to examine a distinct question.

<sup>2</sup> The formal definition of the mean square distance, MSD, is presented in Sect. 2.



**Fig. 1** Relationship between equal choice, equilibrium and average choices in the past  $2 \times 2$  experiments

generated from the play of a small number of games with different procedures used by different researchers. The different procedures make it impossible to address whether the differences between observed and predicted behavior generalize beyond the particular games studied, or whether the deviations observed in each experiment are specific to features of the experimental procedures. Given that the existing evidence suggests that the accuracy of the aggregated proportion of play (across time and player pairs) varies with the distance between the equilibrium and equal choice of actions, we will focus on the aggregated proportion of play across players and repetition of play. Although the literature has explored all levels of aggregation (and disaggregation), we focus on the most aggregated level to test the game level effects.

To make direct comparisons feasible, this paper documents play over a class of constant sum games that are two-person, two-action and have a unique nontrivial mixed strategy equilibrium, and are identical in every respect other than the payoff parameters. Even in this minimalist class of games, choosing the parameters of each game involves trade-offs. One method is to choose parameters so that on a well-defined metric (such as over the payoffs or over the distance from the equilibrium to equal choice), the games are evenly spread out over the entire range. An advantage of this approach is that it could provide the best chance to detect discontinuities in the relationship between the space chosen and behavior if discontinuities exist and enough games are included to detect the discontinuities. A disadvantage of this approach is that the researchers nonetheless choose the specific games and metrics which could in itself bias the games chosen. Alternatively, to avoid potential experimenter bias in choice, we randomly sample from all the possible games rather than choose specific games.<sup>3</sup> Even in this class of games, we will see that the

<sup>3</sup> A risk with randomly sampling games is that the researcher might have bad luck, in that the resulting games could be clustered closely on some dimension in which variation would be critical. This concern

deviation between minimax play and equal probability play of each action systematically generates enough variation to observe on which games the minimax prediction is more or less accurate.

Consistent with the evidence from different experimenters using different procedures and different parameters, we find that behavior is closer to minimax play in games in which the frequency of play predicted by minimax is closer to the equal probability mixture. Indeed, when we use the identical procedures we find a significantly greater relationship than the extant literature finds using different procedures, indicating that the different methods applied by different researchers introduced noise that results in underestimating the strength of an important relationship explaining the accuracy of the minimax hypothesis.

## 2 Experimental design and equilibrium

### 2.1 The random sample of games

Each subject played 500 rounds of one of ten two-player, constant sum games against a fixed, anonymous opponent. Table 1 shows the games. The numbers in each matrix represent probabilities that the row player wins a fixed amount  $w$  on each round. For example, in game 1 if both players choose A, then the row (column) player wins  $w$  with probability 73% (27%). For each round, a player who does not win  $w$  earns 0. In each game,  $w$  was fixed at \$0.04 for games played in the USA and approximately \$0.03 for games played in Israel. A player's payoff for participation in the experiment equaled his payoffs over the 500 rounds of play plus a fixed show-up fee. All transactions were conducted anonymously via networked computers.

All games were randomly chosen as follows. Four numbers,  $p_1$ ,  $p_2$ ,  $p_3$  and  $p_4$ , were drawn from the uniform distribution on the values  $[0.00, 0.01, \dots, 0.99, 1.00]$  to produce a potential game, where  $p_1$ ,  $p_2$ ,  $p_3$  and  $p_4$  were the row player's probability to win  $w$  and  $1 - p_1$ ,  $1 - p_2$ ,  $1 - p_3$  and  $1 - p_4$  were the column player's probabilities to win  $w$ :

Random game	Column player's choice	
	A	B
Row player's choice		
A	$p_1$	$p_2$
B	$p_3$	$p_4$

Footnote 3 continued

can be reduced by increasing the sample of games chosen. Another risk with randomly choosing games is that the researcher needs to draw the games in a credible manner to avoid concern that they somehow did not introduce bias (e.g., picking and choosing which draws to use). To alleviate this concern, researchers could conduct the random draws in a public forum.

**Table 1** Ten random games

Game	Random games			Minimax prediction				Distance between minimax prediction and equal choice mixture
		A	B		A	B		
1	A	0.73	0.60	A	0.275	0.331	<b>0.606</b>	0.12
	B	0.55	0.75	B	0.179	0.215	0.394	
2	A	0.76	0.40	A	0.321	0.333	<b>0.654</b>	0.15
	B	0.23	0.91	B	0.170	0.176	0.346	
3	A	0.93	0.40	A	0.113	0.414	<b>0.527</b>	0.29
	B	0.05	0.64	B	0.101	0.372	0.473	
4	A	0.95	0.67	A	0.159	0.508	<b>0.667</b>	0.31
	B	0.31	0.87	B	0.079	0.254	0.333	
5	A	0.65	0.52	A	0.410	0.431	<b>0.841</b>	0.34
	B	0.23	0.92	B	0.077	0.081	0.159	
6	A	0.62	0.45	A	0.170	0.568	<b>0.738</b>	0.36
	B	0.12	0.60	B	0.060	0.201	0.262	
7	A	0.54	0.46	A	0.269	0.557	<b>0.826</b>	0.37
	B	0.23	0.61	B	0.057	0.117	0.174	
8	A	0.92	0.82	A	0.119	0.664	<b>0.783</b>	0.45
	B	0.53	0.89	B	0.003	0.184	0.217	
9	A	0.92	0.83	A	0.197	0.676	<b>0.873</b>	0.46
	B	0.37	0.99	B	0.029	0.098	0.127	
10	A	0.74	0.73	A	0.203	0.782	<b>0.985</b>	0.57
	B	0.20	0.87	B	0.003	0.012	0.015	
					<b>0.206</b>	0.794		

Cell entries in games show the probability the row player wins

Bold cells indicate the best response to the opponent playing an equal choice mixture

Such a game either has a (weakly) dominant strategy for at least one of the players, or a unique mixed strategy equilibrium in which both players play each of their strategies with positive probability. The first ten games generated in this way that had a unique mixed strategy equilibrium were included in our study.

The current paper examines the way people play these games under two distinct information conditions: high and low information.<sup>4</sup> We examine both high and low-information games to reflect the extant literature that has examined minimax using a range of information conditions from the high-information condition we examine here (i.e., full information) to even less information than in our low-information game. While our focus is on the high-information treatment, we present the low-information treatment results for completeness and robustness.

In the high-information treatment, each player knew the probabilities  $p_1$ ,  $p_2$ ,  $p_3$  and  $p_4$  prior to playing the first round. After each period of play, each player learned what action the other player had chosen. Players were also told whether they received the payoff  $w$  or  $0$ , but were not told whether the other player received  $w$  or  $0$ .<sup>5</sup> Three subject pairs played each game in Table 1 in the high-information condition at Harvard, the Technion (Israel) and the University of Pittsburgh for a total of nine subject pairs per game and 90 total subject pairs across the ten games.

In the low-information treatment, each player knew the structure of the game (i.e., each player knew there were values  $p_1$ – $p_4$  that would determine his lottery payoff and that he was playing against the same opponent every round). However, players did not know the specific realization of  $p_1$ – $p_4$  nor were players ever informed of their opponent's choices. As in the high-information condition, three subject pairs played each of the ten games in Table 1 in the low-information condition at Harvard, the Technion and the University of Pittsburgh for a total of 90 subject pairs across the ten games. We thus examine the behavior of 360 subjects.

## 2.2 Equilibrium

The games in our experiment have two players and two choices per player, and are randomly chosen from the universe of games having a unique equilibrium in nontrivial mixed strategies. In equilibrium, each player chooses his actions with probabilities such that, given the strategy of the other player, no change in probabilities would increase his expected payoff (von Neumann 1928; Nash 1950). Mixed strategy equilibrium predictions are stated here assuming players are expected utility maximizers and the payoffs are expected utilities. Because our games have binary lottery payoffs, the equilibrium predictions can be determined without estimating any unobservable parameters (involving risk aversion) (Kagel and Roth 1995; Roth and Malouf 1979; Wooders and Shachat 2001).<sup>6</sup> This

---

<sup>4</sup> The high-information data were also used in Erev et al. (2007), and the low-information data were also used in Erev et al. (2002). These studies focused on testing learning theories and the development of a statistical tool to facilitate model comparison.

<sup>5</sup> Lottery outcomes were determined independently for each player.

<sup>6</sup> While the binary lottery payoffs allow the minimax predictions to be unambiguously computed, the question of how binary lotteries affect the behavior of the players (who may not be expected utility maximizers) is a different matter. Bereby-Meyer and Roth (2006) showed that players learn more slowly in a repeated prisoner's dilemma game with binary lottery payoffs than in the repeated game with deterministic payoffs with the same expected value.

equilibrium prediction is the game theoretic prediction for the games played under complete information (i.e., in the high-information condition).<sup>7</sup>

### 2.3 Coding the games

Since the games are randomly drawn, there is no a priori meaning to ordering the games. However, based on the past evidence (see Fig. 1), we code the games based on the relationship between the closeness of minimax play and equal probability choice of actions as follows:

#### 2.3.1 Coding games

Let  $MSD_g = \{[0.5 - eq_g(\text{row}, A)]^2 + [0.5 - eq_g(\text{column}, A)]^2\}^{1/2}$ , where  $eq_g(r, X)$  is the mixed strategy equilibrium probability that player  $r$  ( $r = \text{row}, \text{column}$ ) in game  $g$  plays action  $X$  ( $X = A, B$ ). We order the games on the basis of the mean squared distance (MSD) between the equilibrium probabilities and equal choice of actions:  $MSD_1 < MSD_2 < \dots < MSD_{10}$ ; we assign higher numbers to games in which equilibrium is further from the equal choice mixture (0.5, 0.5).

The last column in Table 1 provides the MSD between minimax play and equal choice. For example, in game 10, minimax play indicates row and column players will play A with probability 98.5 and 20.6%, respectively, and row and column players will play AA, AB, BA and BB with probability 20.3, 0.3, 78.2 and 1.2%, respectively. The MSD between minimax play and equal choice mixture is 0.57.

## 3 Results

To study inexperienced, all and experienced play, we examine play during the first 100, all 500 and last 100 rounds, respectively. We first test whether the choice distributions predicted by minimax play describe behavior. We then compare the relative accuracy of minimax across the games.

Tables S2A and S2B (in the online supplemental material) show the proportion of choices in the high- and low-information games, averaging across subjects within each role. For each game, we show the proportion of time each joint choice was made. For example, during the first 100 rounds of high-information game 10, row and column players chose AA, AB, BA and BB 52.0, 27.8, 11.3 and 8.9% of the time, respectively. The proportions of choices each role player made are shown in the margins; row and column players chose A 79.8 and 63.3% of the time, respectively, in high-information game 10.

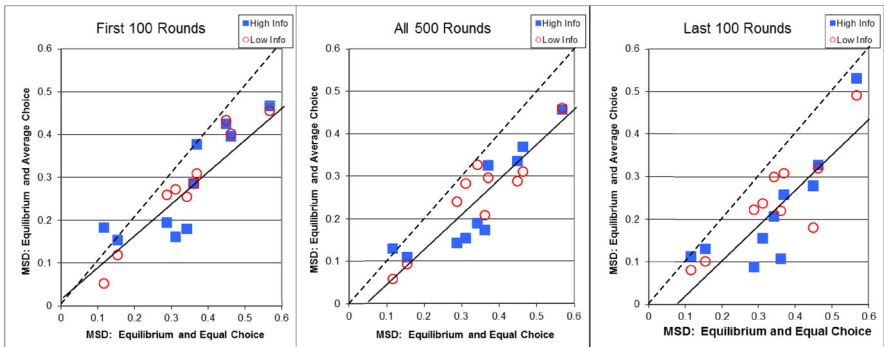
To examine the joint choice distributions predicted by minimax play for each game, we use a  $\chi^2$  goodness-of-fit test to compare minimax play (Table 1) with subject choices (Table S2A and S2B). We conduct the test separately for each experience level, information condition and game. Within each experience level, information condition and game, we aggregate over the nine pairs. For all 60 tests

<sup>7</sup> Since the games are constant sums, many concerns expressed, e.g., by Ochs and Roth (1989) or Bolton and Ockenfels (2000) about experimental control of other aspects of players' preferences are ameliorated.

(three levels of experience times two levels of information times ten games), the minimax prediction is rejected at any conventional level of significance ( $p < 0.01$ ). We repeat this test for each subject pair. Disaggregating the data to each pair, in the high-information games minimax is rejected at the 5% significance level for all 90 subject pairs for every level of experience. In the low-information games, we reject minimax play at the 5% level for 84 of the 90 subject pairs.

**Conclusion 1** Minimax play is rejected based on its ability to predict the exact joint choice distribution of play.

Although we can reject the point predictions for each game, this does not mean that the minimax point prediction fails by the same distance across all games. Figure 2 shows the mean square distance between minimax play and subject behavior as a function of the mean square distance between minimax play and the equal choice mixture. The figure shows results for the first 100 rounds, all 500 rounds and the last 100 rounds. Figure 2 shows that as minimax play is further from the equal choice mixture across games, minimax play is also further from subject behavior. The thick lines in Fig. 2 show the estimated slope from OLS regressions (using one observation per game, aggregating over all subject pairs within each game, resulting in  $N = 10$  observations in each information condition). The estimate indicates a significant ( $p < 0.01$ ) positive relationship between (1) the distance between minimax play and subject choices and (2) the distance between minimax play and equal choice for the first 100, all 500 and the last 100 rounds. Over all 500 rounds, the slope of the OLS line is 0.85 (standard error equals 0.09) and we cannot reject ( $p > 0.10$ ) that the slope is statistically different from 1 (the



Regression:  $MSD(\text{Equilibrium and Average Observed Choice}) = a + b \cdot MSD(\text{Equilibrium and Equal Choice Mixture})$   
 Graphs show the regression line (Solid Line) for the High Information condition

	Parameter Estimates					
	First 100 Rounds		All 500 Rounds		Last 100 Rounds	
	a	b	a	b	a	b
High Information	0.02	0.76	-0.03	0.79	-0.06	0.82
Low Info	-0.03	0.92	-0.01	0.78	-0.01	0.74
p-value	>.25	<.01	>.25	<.01	>.25	<.01
N	10		10		10	

For all regressions, a is always insignificant ( $p > 0.25$ ); For all regressions, b is always significant ( $p < 0.01$ )

**Fig. 2** Relationship between equal choice, equilibrium and average choices (dashed line 45 degree line; solid line estimate from the OLS regression reported below)



dashed line in all of the figures). The regressions also show that there is no significant difference in the estimated slope for the low- and high-information games ( $p > 0.10$ ).

**Conclusion 2** Minimax play is further from subject choices, the further it is from the equal choice mixture.

An important motivation for examining play across a sample of games, holding all else constant (e.g., including the subject population, the stakes, and experimental procedures and instructions) is that we can isolate the relationship between games, *ceteris paribus*. Figure 1 shows that the slope of the relationship between the MSD of minimax and equal choice and the MSD of average subject choices and minimax has a value of 0.55, whereas Fig. 2 shows that this slope is 0.85. OLS regressions combining the data from the current study (using all 500 rounds) and the data from the literature and using one observation per game (again, averaging over all pairs and all rounds of play) indicates that the positive relationship is significantly larger in the games using the identical design collected in this study compared to the games from the literature using different designs ( $p < 0.001$ ). In other words, collecting data across a space of games, *ceteris paribus*, eliminated noise due to the different designs of the past researchers to more clearly observe the relationship across the games.

Figure 1 also suggests that there may be a non-linear relationship between the MSD of minimax and equal choice and the MSD of average subject choices and minimax. To examine the possibility of this non-linear relationship, we re-estimated the OLS regressions adding the squared term of the MSD of minimax and equal choice. With the games using different designs from the literature, this non-linear term is highly significant ( $p < 0.01$ ). Similarly, as seen in Fig. 2, using the data from the games with the identical design, the non-linear squared term is marginally significant ( $p = 0.056$ ) in the high-information games using all 500 rounds, but is not significant using the low-information games ( $p > 0.25$ ).

Finally, note that the OLS regressions make strong assumptions relying on asymptotic approximations that are questionable given just 13 observations with the games with different designs (from the literature) and with our current ten games with the identical design. To address this concern, we examined the Spearman rank order correlation across the games. In particular, we ranked the games from the smallest to largest MSD between minimax and equal choice separately for the 13 games from the literature and the ten games we ran (see Table 1). We grouped games from the literature with the identical MSD between minimax and equal choice into a single rank and took the average MSD between subject's choices and minimax for the subject's behavior for these ranks. This results in eight and ten distinct ranks for the games with different designs and for our games with the identical design, respectively. The Spearman rank order correlation for the games with different designs from the literature is 0.90. The Spearman rank order correlation for the high- and low-information conditions with the same design is 0.98 and 0.95, respectively, and all correlations are highly significant ( $p < 0.001$ ). These non-parametric correlation tests support Conclusion

2 that subjects choices are further from minimax the further minimax is from equal choice.<sup>8</sup>

## 4 Discussion

Our experimental design lets us make inferences about play over a class of games. We find across all levels of experience that minimax play is unable to describe the joint choice distribution across the population of games. However, minimax play is closer to subject average choices in games in which minimax play is closer to the equal choice mixture. The latter result could not have been observed without examining many games in the sample space.

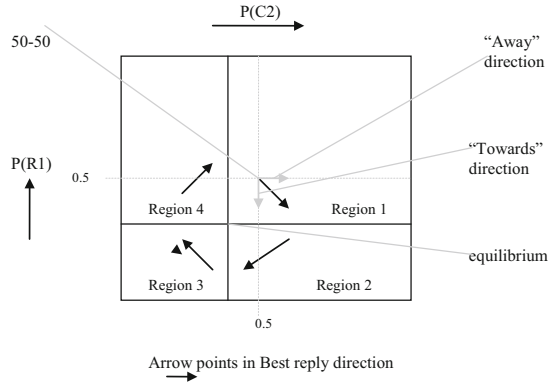
Our core result helps explain the relative accuracy and inaccuracy of the minimax hypotheses to describe behavior across many of the past studies using a two-player, two-action design to test the minimax hypothesis.<sup>9</sup> Despite often dramatic differences in designs that past researchers have used, Fig. 1 shows that across these different designs play is closer to subject average choices in games in which minimax play is closer to the random equal choice mixture. Thus, one reason the behavior in some of these past studies is closer to the minimax hypotheses than some of the other studies is due to the specific parameters of the games in which the minimax prediction is closer to equal choice. While the current results do not show that other aspects of the design are unimportant for explaining the observed accuracy or inaccuracy of the minimax hypothesis, they do indicate that the location of the predicted play relative to equal choice will significantly affect the hypothesis' accuracy.

Our results provide an empirical basis for advances in understanding the descriptive accuracy of minimax play. Examining a class of games, we documented a critical characteristic that systematically affects behavior. Expanding to broader spaces (for instance across more players or actions) will further indicate whether the distance from equal choice of actions is robust to higher dimensions. Continuing research in this direction will move toward an understanding of what is necessary for a descriptive model to account for behavior unconditional on the location within

---

<sup>8</sup> We also compared the difference in the MSD of subject choice and minimax between all adjacent ranks (i.e., rank 1 vs. rank 2, rank 2 vs. rank 3, etc.). For the seven comparisons in the games with different designs, in six the MSD of subject choices and minimax is greater for the higher ranked games, and for all nine comparisons of the games with the identical design we find that the MSD of the subject choices and minimax is greater for the higher ranked games (using the average of the high and low-information games). If the null hypothesis is no correlation such that there is a 50 percent chance of a greater difference in the MSD of subject choices and minimax across the adjacent ranks, then these results reject at the  $p = 0.06$  and  $p < 0.01$  level no correlation for the games with different and the identical design, respectively, in favor of a positive correlation in which subject's choices are further from minimax the further minimax is from equal choice.

<sup>9</sup> This paper presented evidence on choice frequency predictions of the minimax hypothesis. However, minimax makes predictions on other aspects of play, such as serial correlation of choices. In an earlier draft of this paper, we explored other predictions of the minimax hypothesis and found that other aspects of the minimax hypothesis vary not only by the game, but also by asymmetries in the roles between the row and column players.

**Fig. 3** Best reply dynamics

the sample space, and for what behaviors are idiosyncratic to the location in the sample space.

Finally, while the focus of this paper is methodological, there are several behavioral effects that can contribute to our core result stated in conclusion 2 and depicted in Figs. 1 and 2. First, there is a tendency for some people to exhibit a 'level-0' or a 'level-1' behavior (Stahl 1996). Level-0 behavior involves random choice and level-1 involves the best reply to the belief that other players are level-0. When the minimax prediction is near equal choice, both level-0 and level-1 choices approximate the minimax prediction. Second, there is evidence suggesting a slow learning process that reflects the joint effect of exploration and an attempt to select the option that has led to the best outcomes in previous trials. In the current context, processes of this type imply quick convergence when minimax predictions are near equal choice (see Ochs 1995; Erev and Roth 1998; Feltovich 2000). Similar predictions are also derived under the assumption that the players rely on small samples of past experiences (Erev and Roth 2014). Third, note that there is robust experimental evidence that players in games with mixed strategy equilibria tend initially to use less extreme mixed strategies than the equilibrium predicts, and that with experience they tend to move slowly in the direction of best response (e.g., Ochs 1995; Erev and Roth 1998; Feltovich 2000). If both players shade their mixed strategies toward the equal choice initially, it is easy to show that one player's best response will be to increase the action that the mixed strategy equilibrium predicts he will play more often (call this the toward player) and the other player's best response will be to decrease the action that the mixed strategy equilibrium predicts she will play more often (call this the away player). Figure 3 provides an example of this situation with arrows to represent the best response dynamics. Until the toward player's mixed strategy behavior involves using the action minimax predicts he will use more often with equal or higher probability than minimax predicts, the away player's best response will continue to move further away from her mixed strategy equilibrium, and the MSD between the players actions and minimax will increase (assuming the players are adjusting their mixed strategies at similar rates). Following Erev and Roth (1998) who observed very slow learning in the best response direction in games with a unique mixed strategy equilibrium, we would

conjecture that for games where the mixed strategy equilibrium is farther from equal choice, it will take longer for (slow) learners to adjust their mixed strategies, and thus larger deviations will occur and persist in these games.<sup>10</sup> The evidence presented in Figs. 1 and 2 is consistent with this conjecture, and subsequent research can attempt to disentangle these effects and examine whether alternative hypotheses may also explain the observed behavior presented here.

## References

- Bereby-Meyer, Yoella, & Roth, Alvin. (2006). The speed of learning in noisy games: Partial reinforcement and the sustainability of cooperation. *The American Economic Review*, 96(4), 1029–1042.
- Bolton, Gary, & Ockenfels, Axel. (2000). ERC: A theory of equity, reciprocity and competition. *The American Economic Review*, 90(1), 166–193.
- Choi, Raymond, Gale, Douglas, & Kariv, Shachar. (2007). Consistency and heterogeneity of individual behavior under uncertainty. *American Economic Review*, 97, 1921–1938.
- Erev, Ido, Ert, Eyal, & Roth, Alvin. (2010a). A choice prediction competition for market entry games: An introduction. *Games*, 1, 117–136.
- Erev, Ido, Ert, Eyal, Roth, Alvin, Haruvy, Ernan, Herzog, Stefan, Hau, Robin, et al. (2010b). A choice prediction competition, for choices from experience and from description. *Journal of Behavioral Decision Making*, 23, 15–47.
- Erev, Ido, & Roth, Alvin. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4), 848–881.
- Erev, Ido, & Roth, Alvin. (2014). Maximization, learning and economic behavior. *Proceedings of the National Academy of Sciences*, 111(Supplement 3), 10818–10825.
- Erev, Ido, Roth, Alvin, Slonim, Robert, & Barron, Greg. (2002). Predictive value and the usefulness of game theoretic models. *International Journal of Forecasting*, 18(3), 359–368.
- Erev, Ido, Roth, Alvin, Slonim, Robert, & Barron, Greg. (2007). Learning and equilibrium as useful approximations: Accuracy of prediction on randomly selected constant sum games. *Economic Theory Special Issue Symposium on Behavioral Game Theory*, 33(1), 29–51.
- Feltovich, Nick. (2000). Reinforcement-based vs. beliefs-based learning in experimental asymmetric-information games. *Econometrica*, 68, 605–641.
- Kagel, John, & Roth, Alvin. (1995). *Handbook of Experimental Economics*. Princeton: Princeton University Press.
- Malcolm, David and Bernhardt Lieberman, 1965. “The behavior of responsive individuals playing a two-person, zero-sum game requiring the use of mixed strategies,” *Psychonomic Science*, 2(12), 373–374.
- Mookherjee, Dilip, & Sopher, Barry. (1994). Learning behavior in an experimental matching pennies game. *Games and Economic Behavior*, 7, 62–91.
- Nash, John. (1950). Equilibrium points in  $n$ -person games. *Proceedings of the National Academy of Science*, 36, 48–49.
- Neumann, J. V. (1928). Zur theorie der Gesellschaftespiele. *Mathematische Annalen*, 100, 295–320.
- Ochs, Jack. (1995). Simple games with unique mixed strategy equilibrium: An experimental study. *Games and Economic Behavior*, 10, 202–217.
- Ochs, Jack, & Roth, Alvin. (1989). An experimental study sequential bargaining. *The American Economic Review*, 79(3), 355–384.
- O’Neill, Barry. (1987). Nonmetric test of the minimax theory of two-person zerosum games. *Proceedings of the National Academy of Sciences*, 84, 2106–2109.

<sup>10</sup> Shachat and Swarthout (2004) describe observing a similar dynamic in which the toward player’s mixed strategy play moves to the minimax play, and then the other player is more likely to play both actions equally. This behavior would lead to our results as well; if one player (the toward player) has converged to minimax behavior, and the other player (the away player) is near equal choice, then the MSD of this behavior will be larger the further the players’ mixed strategy is away from equal choice.

- Rosenthal, Robert, Shachat, Jason, & Walker, Mark. (2003). Hide and Seek in Arizona. *International Journal of Game Theory*, 32, 273–293.
- Roth, Alvin, & Malouf, Michael. (1979). Game-theoretic models and the role of information in bargaining. *Psychological Review*, 86, 574–594.
- Roth, A., Prasnikar, V., Okuno-Fujiwara, M., & Zamir, S. (1991). Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. *American Economic Review*, 81, 1068–1095.
- Shachat, Jason, & Swarthout, Todd. (2004). Do we detect and exploit mixed strategy play by opponents? *Mathematical Methods of Operations Research*, 59(3), 359–373.
- Slonim, Robert, & Roth, Alvin. (1998). Learning in high stakes ultimatum games: An experiment in the Slovak Republic. *Econometrica*, 66(3), 569–596.
- Stahl, Dale. (1996). Boundedly rational rule learning in a guessing game. *Games and Economic Behavior*, 16, 303–330.
- Suppes, Patrick, & Atkinson, Richard. (1960). *Markov learning models for multiperson interactions*. Stanford: Stanford University Press.
- Wooders, John, & Shachat, Jason. (2001). On the Irrelevance of risk attitudes in repeated two-outcome games. *Games and Economic Behavior*, 34, 11–33.