

What Happens in the Field Stays in the Field: Exploring Whether Professionals Play Minimax in Laboratory Experiments

Steven D. Levitt, John A. List, and David H. Reiley*

First version: August 2007

This draft: July 2009

Abstract

The minimax argument represents game theory in its most elegant form: simple but with stark predictions. Although some of these predictions have been met with reasonable success in the field, experimental data have generally not provided results close to the theoretical predictions. In a striking study, Palacios-Huerta and Volij (2007) present evidence that potentially resolves this puzzle: both amateur and professional soccer players play nearly *exact* minimax strategies in laboratory experiments. In this paper, we establish important bounds on these results by examining the behavior of four distinct subject pools: college students, bridge professionals, world-class poker players, who have vast experience with high-stakes randomization in card games, and American professional soccer players. In contrast to Palacios-Huerta and Volij's results, we find little evidence that real-world experience transfers to the lab in these games—indeed, similar to previous experimental results, all four subject pools provide choices that are generally not close to minimax predictions. We use two additional pieces of evidence to explore why professionals do not perform well in the lab: (1) complementary experimental treatments that pit professionals against preprogrammed computers, and (2) post-experiment questionnaires. The most likely explanation is that these professionals are unable to transfer their skills at randomization from the familiar context of the field to the unfamiliar context of the lab.

JEL: C72 (Noncooperative games); C9 (Design of Experiments)

Key words: Mixed Strategy, Minimax; Laboratory Experiments

* We would like to thank Editor David Levine and three anonymous referees for valuable comments. Ignacio Palacios-Huerta and Jesse Shapiro provided helpful conversations. Phil Gordon was instrumental in our efforts to recruit world-class poker players. Omar Al-Ubaydli, David Caballero, Dwyer Gunn, Bill Hessert, Ryan Johnson, Min Lee, Randall Lewis, Andrew Sherman, Alec Smith, Brittany Smith, Dean Strachan, and especially Lisandra Rickards provided fantastic research assistance. Author affiliations: Levitt and List: Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago, IL 60637. Reiley: Department of Economics, University of Arizona, 401 McClelland Hall, Tucson, Arizona 85721. E-mail addresses: slevitt@uchicago.edu, jlist@uchicago.edu, reiley@eller.arizona.edu

I. Introduction

John von Neumann's (1928) Minimax theorem preceded Nash Equilibrium as the first general framework for understanding of play in strategic situations. The underlying logic of the minimax argument has subsequently been applied broadly—from models of firm, animal, and plant competition to the optimal actions of nations at war. In zero-sum games with unique mixed-strategy equilibria, minimax logic has an intuitive appeal: one needs to randomize strategies in order to prevent exploitation by one's opponent. A nagging issue is that subjects in laboratory studies typically do not play near the predictions of minimax (see, e.g., Lieberman, 1960; 1962; Brayer, 1964; Messick, 1967; Fox, 1972; Brown and Rosenthal, 1990; Rosenthal et al., 2003). Perhaps this finding should not have come as a surprise given that experimental subjects are unable to produce a random series of responses, even when directed to do so (Budescu and Rapoport, 1992). O'Neill (1991, p. 506) aptly summarizes these results by noting that "by the mid-1960s, non-cooperative theory had received so little support that laboratory tests ceased almost completely."

Recent evidence from field data has provided a renewed sense of optimism, however. Walker and Wooders (2001) analyze serve choices in Grand Slam tennis matches and report similar win rates across various strategies, a result consistent with the minimax equilibrium prediction. Yet, they find that the players switch from one strategy to another too often, a result at odds with minimax theory but consonant with laboratory experimental evidence. Hsu et al. (2007) analyze a broader tennis data set, finding results even more consistent with theory: not only are win rates similar across strategies, but individual play is serially independent. Complementing these data are results from Chiappori et al. (2002) and Palacios-Huerta (2003), who examine penalty kicks in

professional soccer games. Both studies report that winning probabilities are identical across strategies and that choices are serially independent.

Combined, these two strands of literature present an important puzzle: why do controlled laboratory tests of minimax systematically provide data far from minimax predictions, whereas less controlled tests using field data appear to confirm theory? Perhaps the tests using field data lack statistical power to reject minimax play. Or, maybe the laboratory has not provided the appropriate environment—for example, crucial experience and context—for subjects to learn the gaming rules to produce equilibrium play.

Palacios-Huerta and Volij (2007) provide a third possible resolution to this puzzle: typical subjects in laboratory experiments do not have uniformly high skill at playing games with mixed-strategy equilibria, but the subset of individuals who excel at tennis or penalty kicks do have uniformly high skill, and they are able to transfer their skills from one setting to another. In support of this conjecture, using two standard zero-sum laboratory games with both amateur and professional soccer players as experimental subjects, they report striking evidence that both subject pools use exact minimax strategies. To the best of our knowledge, this study is the first laboratory experiment to show that subjects can both i) play strategies in the predicted equilibrium proportions, and ii) generate a sequence of choices that are serially independent. The authors attribute their result to skills learned in soccer that are subsequently transferred to laboratory card games.

In an effort to understand when field behavior does and does not translate into lab performance, this paper begins by summarizing data from three distinct subject pools — undergraduate students, professional bridge players, and professional poker players — playing the same two zero-sum laboratory games as in Palacios-Huerta and Volij (2007).

Members of both of our professional subject pools have extensive experience thinking analytically in card games. For our purposes, however, there is one crucial difference between bridge and poker players: there is virtually no role for mixed strategies in bridge. In contrast, randomization is an integral component of skillful poker, as noted, for example, by Friedman (1971, B-764): “it is quite clear to those who have played much poker that some sort of mixed strategy...must be used.”¹

Empirical results from all three subject pools parallel those previously reported in the literature using laboratory experiments with student subjects: play is shown to deviate considerably from the theoretical predictions (see, e.g., Lieberman, 1960; 1962; Brayer, 1964; Messick, 1967; Fox, 1972; Brown and Rosenthal, 1990; Rosenthal et al., 2003; Palacios-Huerta and Volij, 2007). Importantly, professional poker players play no closer to minimax than students and bridge professionals, and far from minimax predictions. This finding holds when the professionals compete against other players, as well as when they are informed that they are playing against a computer preprogrammed to exploit individual deviations from optimal play.

These results induced us to collect data from our own sample of professional soccer players drawn from three Major League Soccer (MLS) teams: the Los Angeles Galaxy, Chivas USA, and Real Salt Lake. Again, the empirical results mirror those found with other subject pools, both in our own experiments and in the previous literature: play is much further from equilibrium than observed in Palacios-Huerta and Volij (2007). In light of the large body of psychological evidence that reports limited transfer of learning across

¹ Hirschberg et al. (2008) provide further anecdotal evidence of the importance of randomization and empirical documentation of poker players randomizing. Using thousands of observations, they find that online poker players equalize payoffs across strategies when mixing. That paper also shows that mixing is an extremely common occurrence in the game examined (heads-up limit hold 'em).

tasks (Loewenstein, 1999), we suspect that our failure to find that play in the lab closely approximates minimax predictions is due to the fact that the two zero-sum games themselves are not ideal representations of what the subjects actually face in the field, or at least the players are not recognizing them as such. To dig a level deeper into this hypothesis, we conducted a post-experimental survey inquiring how the professional soccer players interpreted the experimental game. Consonant with the data patterns observed, not one soccer player who participated in the experiment spontaneously responded that the experiment reminded him of penalty kicks. Even when specifically prompted with a question about penalty kicks, many of the subjects saw little connection between the lab game and penalty kicks.

II. Experimental Design

We chose to follow Palacios-Huerta and Volij (2007) in using two different matrix games, which we described to subjects as “Hide and Seek” and “Four-Card Barry.” Hide and Seek is a 2x2 matrix game taken from Rosenthal et al. (2003); Four-Card Barry is a 4x4 matrix game developed by O’Neill (1987).² As Figures 1a and 1b demonstrate, both games are two player zero-sum games with non-uniform equilibrium mixing proportions.

A. Student Subjects

To provide a baseline of comparison, our examination begins with an exploration of undergraduate student play in these two games. We included a total of forty-six students from the University of Arizona—twenty-two students participated as subjects in two

² The O’Neill experiment was seminal in the sense that it moved experimental tests of minimax theory to an environment that required fewer assumptions on players’ utility functions. Prior to O’Neill (1987), previous experimenters assumed that utility depended only on the players’ own payoff and furthermore was a linear function of that payoff. By restricting the game to two outcomes—win or lose the same dollar amount—O’Neill was able to construct a matrix with the property that a players’ minimax strategy is invariant over reasonable utility functions. Mark Walker and John Wooders suggested the name “Four-Card Barry” to us.

sessions of Hide and Seek and twenty-four participated as subjects in three sessions of Four-Card Barry. No subject participated in more than one session.

Following Rosenthal et al. (2003), in the Hide and Seek treatment we had each pair of participants sit opposite each other, with a game conductor sitting at the side of the table. The conductor gave each participant the instructions for Hide and Seek (see the Supplementary Appendix) and read them aloud. Then the participants played several practice rounds until each was sure she was ready to play for real money. Participants made their decisions by playing either a red card or a black card,³ with both players' decisions revealed simultaneously. As specified in the instructions, the game conductor sometimes rolled a six-sided die to determine the winner of a round. Participants played a total of 150 rounds, switching roles after 75 rounds. Each round, one player won a payoff of \$0.25. Each session of Hide and Seek lasted less than an hour. The payoff matrix in our version of Hide and Seek follows that of Rosenthal et al. (2003), and therefore differs slightly from the 2x2 penalty-kick game studied by Palacios-Huerta and Volij (2007).⁴

As in Hide and Seek, we implemented Four-Card Barry in a manner that closely followed the literature. Each pair of participants for Four-Card Barry sat opposite each other, with a conductor sitting at the side of the table. The conductor gave each participant the instructions for Four-Card Barry (see the Supplementary Appendix) and read them aloud. The participants then played as many practice rounds as they wished until both were

³ We used regular playing cards, typically with a ten of a red suit and a ten of a black suit for each player.

⁴ Palacios-Huerta and Volij (2007) created a payoff matrix based on the empirical success percentages from professional penalty kicks. Since we intended to play our game with multiple subject pools, we elected instead to implement a game that was easy for the subjects to understand, but was asymmetric and had equilibrium mixing proportions different from 50:50 for each player. Consequently, the equilibrium mixing proportions differ slightly from those of Palacios-Huerta and Volij (2007). Palacios-Huerta and Volij (2007) have equilibrium mixing proportions of approximately 36:64 for Player 1 and 55:45 for Player 2; Hide and Seek has equilibrium mixing proportions of 67:33 for each player. If anything, this change provides theory with a better chance to succeed, as a two-thirds mixing proportion might be cognitively easier to execute than a more complicated proportion like 36:64 or 55:45.

ready to play for real money. Participants made their decisions by playing one of their four cards, with both players' decisions revealed simultaneously. Participants played a total of 150 rounds, switching roles after the first 75 rounds. Each round produced a winner who received a \$0.25 payoff; each session of Four-Card Barry typically lasted less than an hour.

This variant of Four-Card Barry is identical to Palacios-Huerta and Volij (2007), with two minor exceptions. First, since we intended to play this game with professional card players, we chose to use regular playing cards with all four of our subject pools. That is, instead of the colored cards (Red, Brown, Purple, Green) used by Palacios-Huerta and Volij (2007), we gave each player one card of each suit (Spade, Heart, Club, Diamond).⁵ Second, to provide additional insights into the ability of subjects to transfer knowledge across tasks, we had the players switch roles halfway through the 150 rounds of the game. Given that previous results show little evidence that play changes over periods (Rosenthal et al., 2003, and Palacios-Huerta and Volij, 2007), this change is likely innocuous.

B. Professional Subjects

We used three types of professionals as subjects in our artefactual field experiment (Harrison and List, 2004).⁶ The first is professional poker players. Our poker treatments were conducted in 2006 at the World Series of Poker, in Las Vegas, NV. There were 130 participants: 44 playing Hide and Seek, 52 playing Four-Card Barry, and 34 who played against a preprogrammed computer (see below). Among our sample of players, the self-reported average number of hours spent per week playing poker is 25. Over 87 percent of

⁵ Each player received all four cards of the same rank, either four nines or four tens. We deliberately avoided using aces in order to avoid having the ace be focal, as most decks of playing cards make the ace of spades much larger than the aces of the other suits.

⁶ For space purposes, we suppress further discussion of our world-class bridge players. As aforementioned, bridge players offer an interesting counterpoint to poker players because, unlike poker, there is virtually no role for mixed strategies in bridge. Bridge players, like our other subject pools described below, deviated systematically from minimax play. Full results on bridge players are available in the Supplementary Appendix and in an earlier version of this paper (Levitt, List, and Reiley 2007).

the players reported to have made money playing poker in the previous year, with average annual earnings of \$120,111. The sample included 11 individuals who had won either a World Series of Poker Bracelet or a World Poker Tour event. Recruiting was done through the distribution of flyers and face-to-face solicitation at the World Series of Poker venue (the Rio Hotel). All treatments were carried out in our hotel suites at the Rio Hotel, which hosted the World Series of Poker. Subjects were paid \$1 per successful play.⁷ The experiment lasted, in most cases, no longer than one hour.

Our second professional subject pool, like Palacios-Huerta and Volij (2007), is professional soccer players. We obtained permission to run experiments with three MLS teams: the Los Angeles Galaxy, Chivas USA, and Real Salt Lake. Each of these clubs granted us access to their locker room for two hours or less. Given the time constraint, we limited our soccer player treatments to the 4X4 O'Neill game, which yielded the most striking results reported in Palacios-Huerta and Volij (2007). We played Four-Card Barry with a total of thirty-two players from the three MLS teams, typically with four or five game conductors simultaneously administering the game to different pairs.

These thirty-two players included thirty roster players, plus one team trainer (who had previously played intercollegiate soccer) and one youth player (a goalkeeper) who trained with the team, but was not yet on the official roster. Five of the thirty-two players were goalkeepers, and following Palacios-Huerta and Volij (2007) we made sure to have all five goalkeepers play against non-goalkeepers. Again, we followed identical protocol to that discussed above with poker players, except that the treatments were carried out in the locker room of the professional soccer clubs.

⁷ In this regard, we followed Fehr and List (2004) and Haigh and List (2005) in using larger payoffs for the professionals than the students. This was done to provide more comparable payoffs on an opportunity-cost scale and maintain the professionals' attention during the games.

C. Human vs. Computer Treatments

With the poker players, we complemented the human-human treatments with two computer treatments.⁸ The first of these involved a computer programmed to play minimax for the first 15 periods and which thereafter mixed between minimax and the action with the highest predicted payoff based on a logit/multinomial logit model that used as right-hand-side variables the previous moves by both the human and the computer.⁹ As more data accumulated, the program gave increasing weight to the model's predictions.¹⁰ Given the nature of what we desired to learn from this exercise, the instructions to this game (included in the Supplementary Appendix) told the player that "...we have programmed the computer to play the theoretically correct strategy in this game. In addition, any deviations that your play has from this correct style of play will be taken advantage of by the computer." These instructions reflect the fact that the computer plays minimax initially, but responds to human play that is far from minimax by attempting to exploit it. As a shorthand, we refer to this treatment as the "optimal computer."

⁸ This line of research originates with Messick (1967), who conducted a three choice, two-player repeated experiment where human subjects played against computer algorithms. See also Fox (1972), Coricelli (2004), Shachat and Swarthout (2004), and Spiliopoulos (2007).

⁹ It is important to build in this learning component because if the computer played Minimax equilibrium probabilities regardless of the response of the human opponent, the human should be indifferent between the choice of actions. By programming the computer to exploit play that is off the equilibrium path, we provided subjects with an incentive to play minimax proportions in order to avoid exploitation.

¹⁰ To be more specific, this involved a two step procedure: first, we estimated a predicted choice for the human player via a logit regression model (for Hide and Seek) or a multinomial logit model (Four-Card Barry) in which the only right-hand-side variables were dummies for (up to) three previous actions by the human and the computer, and calculated a best response to that predicted choice. Second, we averaged that best response together with the theoretical equilibrium ratios (1/3, 2/3 in Hide and Seek; 2/5, 1/5, 1/5, and 1/5 in Four-Card Barry) to provide the computer's mixed-strategy proportions for the next round. We increase the weight given to the logit prediction over time as those predictions are based on more data. In Hide and Seek we used the simple average of the predicted logit best-response strategy and the theoretical equilibrium strategy for periods 16-35; in periods after 35 we used a weight of $\frac{3}{4}$ on our predicted best response and $\frac{1}{4}$ on the equilibrium play. Four-Card Barry was identical except that we used a $\frac{1}{4}$ weight on our predicted best response and $\frac{3}{4}$ weight on the equilibrium ratio for periods 16-25.

Our second computer treatment involved programming the computer to play sub-optimally. In particular, we chose a simple algorithm whereby the computer randomly chose between the available actions with equal probabilities (whereas the optimal mix was 67:33 in the 2x2 game, and 40:20:20:20 in the 4x4 game). This was a static strategy, and no computer learning component was built into this treatment. We refer to this treatment as the “naïve” computer program. The instructions for this treatment differed from those of the first computer treatment only by the omission of the sentence saying that the computer had been programmed to play optimally.

In total, we had 34 participants in the computer treatments. To maximize sample sizes, we allocated the participants so that each would play both games. But we did not vary the computer algorithm condition: if a subject was randomly placed in the “optimal” condition, for example, then she was in that condition for both games. And, to control for order effects, the computer program randomly decided whether Hide and Seek or Four-Card Barry would be played first. In aggregate, 21 subjects played both Hide and Seek or Four-Card Barry against an optimally programmed computer, and 13 played both games against a computer programmed to play an exploitable strategy.

III. Empirical Results

As noted in the literature, if subjects are playing the unique minimax equilibrium, then the data generated should conform to three key predictions: (1) for all players combined, the aggregate marginal and joint distributions of actions should correspond to that predicted by equilibrium play, (2) for each particular pair of players, the marginal and joint distribution of actions should correspond to that predicted by equilibrium play, and (3)

actions should be serially uncorrelated.¹¹ In what follows, we report our results parsed by subject pool.

A. Hide-and-Seek Results

Table 1 summarizes our findings for Hide and Seek. Each column corresponds to a different subject pool. The first two columns provide our data on college students and poker players. For purposes of comparison, we also report the soccer player results obtained by Palacios-Huertas and Volij (2007) in their 2x2 game. The top two rows of the uppermost panel in Table 1 show sample sizes. The lower panels provide tests of the relevant hypotheses described above. Readers interested in greater detail regarding these findings are directed to the Tables in the Supplementary Appendix.

Panel I in Table 1 reports p-values for rejecting the null hypothesis that the *aggregate* frequencies match those of minimax play. The first and second rows show results corresponding to the marginal distributions for those playing pursuer (or seeker) and evader, (or hider).¹² The third row reports results for the joint distribution of play, showing whether combinations of actions (for example, black-black plays) match minimax predictions. For both of our subject groups, all three of these hypotheses are rejected at the $p < .01$ level. And, importantly, the magnitude of the aggregate deviations from equilibrium is substantial. Theory predicts that both players should adopt a 67:33 red to

¹¹ Palacios-Huerta and Volij also test a fourth hypothesis that expected win rates across strategies should be equal to each other and to the predicted equilibrium win rate. This test is important for field studies such as Walker and Wooders (2001) and Palacios-Huerta (2003), where the true payoff matrix is unknown, but is unnecessary in this setting since it is superfluous once hypothesis (2) is tested, in that it follows mechanically from (2). Further, randomization to produce the winner, particularly in the 2x2 game, introduces more noise and lower power for a test of win rates versus a test of choice frequencies. We therefore conserve space and place these results in the Supplementary Appendix; but we should note that for over sixty percent of both students and poker players, we can reject equality of success rates at the $p < .05$ level. And, many of the rejections show behavior far removed from the minimax prediction. In stark contrast, Palacios-Huerta and Volij (2007) cannot reject for a single professional soccer player.

¹² These p-values are obtained from Pearson's Chi-square test for goodness of fit, using one degree of freedom for the test of marginal frequencies and three degrees of freedom for the test of joint frequencies.

black ratio. Students played red 61 percent of the time; poker players only 56 percent. Note, however, that this is also one of the few tests on which the soccer players in Palacios-Huerta and Volij (2007) strayed somewhat from minimax, as shown in columns 3 and 4.

Panel II in Table 1 continues to focus on action frequencies, but differs from the top panel in reporting results for *individual* pairs of players, rather than summarizing the aggregate data. Instead of reporting p-values, in this case we show the fraction of individual players for whom we can reject the null hypothesis that the player's actions match minimax play when acting as the pursuer or evader at the $p < .05$ level.¹³ The third row in panel II reports a similar statistic, but for the joint play by each pair. In contrast to panel I, *large* numbers represent violations of minimax in this part of the table.¹⁴ We also provide the distribution of individual choice frequencies in Figures 2a and 2b, along with the theoretical distribution that these choice frequencies should have under minimax.

For both students and poker players, we find considerable departures from minimax play. Whether in the role of pursuer or evader, more than half of the subjects engage in play that is inconsistent with minimax behavior at the $p < .05$ significance level.¹⁵ We are able to reject the hypothesis that both players are jointly following minimax in at least 75 percent of the pairs. In roughly one-third of the pairs, neither of the players' actions is consistent with minimax. Most of the deviations take the form of playing red too infrequently: nearly one-third of the students play red less than half of the time (minimax predicts red 67 percent of the time); approximately one-fourth of the poker players chose

¹³As before, we use a Pearson Chi-square test with one degree of freedom for the marginal frequencies and three degrees of freedom for the joint frequencies.

¹⁴We also examined the frequency with which *neither* player follows minimax, which represents a stronger test of the theory since if one player is following minimax, both players receive the equilibrium payoffs, regardless of the strategy the other follows. The patterns generally follow those discussed above.

¹⁵For purposes of comparison, when we program computers to naively play a 50:50 mix, we are able to reject the null at the .05 level in 65 percent of the cases.

red less than half of the time. Note that our findings in Panel II differ starkly from Palacios-Huerta and Volij (2007), as revealed in columns 3 and 4. In their sample, rejections were no more frequent than chance would predict under the null.

Finally, panel III of Table 1 presents the percentage of players for whom we can reject the null hypothesis of no serial correlation in actions, based on the runs test of Gibbons and Chakraborti (1992). The students and poker players fare much better on this test than the other tests, with “only” about 20 percent of the players exhibiting play that significantly deviates from the no-serial-correlation null, although this is still worse than Palacios-Huerta and Volij’s (2007) soccer players. When Palacios-Huerta and Volij do reject, particularly with college soccer players, it tends to be for too many runs, which means too frequent switching of strategies.¹⁶ By contrast, our rejections of serial independence are at least as likely to be for too few runs as for too many runs, indicating that players frequently fail to switch often enough.¹⁷

In sum, the results we obtain using either undergraduate students or professional card players parallel those previously reported in the literature (see, e.g., Brown and Rosenthal, 1990; Rosenthal et al., 2003). Consistent with Rosenthal et al. (2003), these results hold whether we examine early periods of play or later periods, suggesting that play is not converging to equilibrium.¹⁸

B. Four-Card-Barry Results

¹⁶ This is consistent with Walker and Wooders’ (2001) professional tennis study.

¹⁷ The serial correlation performance is perhaps better than one would expect based on prior individual-level studies in psychology that lead one to conclude that “producing a random series of responses is difficult, if not impossible task to human [subjects], even when they are explicitly instructed” (Wagenaar, 1971, p. 78). However, it is in line with the intuition that subjects competing in dyadic interactions are more likely to yield serially uncorrelated play than in parallel individual choice settings (Budescu and Rapoport, 1992).

¹⁸ See the Supplementary Appendix for the results split by the first and second halves of the treatment. Across all of our 2x2 treatments, the results for panels I-III are similar for the two halves of play. In all cases, however, the frequency of rejection for serially correlated play is lower as subjects gain experience with the game. The supplementary appendices also highlight the economic significance of these departures.

Table 2 presents results from Four-Card Barry. The structure is similar to Table 1, except that Table 2 contains one additional column corresponding to our sample of professional soccer players. Thus, the first three columns of Table 2 contain our data, and columns 4 and 5 report the results of Palacios-Huerta and Volij (2007).

Panel I presents results on aggregate frequencies.¹⁹ For college students and poker players, the aggregate proportions are closer to theory in Four-Card Barry than in the 2x2 game, although we continue to reject at the $p < .01$ level that the column players have the minimax mixing proportions. The substantive magnitudes of these deviations, however, are small: in all cases the aggregate proportions are within a few percentage points of the predicted values. Our professional soccer players provided the largest deviations we found from theory in the aggregate data, where Row players played 43% diamonds and 17% clubs (instead of 40% and 20%), while Column players played 23% clubs and 17% hearts (instead of 20% each). These deviations are reflected in the especially low p-values for soccer players. We find it surprising that our soccer players deviate more from minimax than do students or poker players, given that the aggregate frequencies of soccer players – even intramural college soccer players – in Palacios-Huerta and Volij (2007) so closely matched the predictions of theory. Indeed, their data are stark in that only one time in ten would such a result be obtained by chance, even if every player was following the minimax strategy.²⁰

Panel II of Table 2 paints a similar picture when analyzing mixing proportions at the individual level. Across our various samples, we are able to reject at $p < .05$ the null

¹⁹ The Chi-squared tests for the 4x4 games use 3 degrees of freedom for the marginal distributions and 15 degrees of freedom for the joint distributions in panels I and II. In Panel III, play is broken down into diamond versus non-diamond plays and the analysis proceeds as in the 2x2 game.

²⁰ Theory predicts mixing proportions of 40:20:20:20. In Palacios-Huerta and Volij (2007), the observed aggregate frequencies among professionals were 39.8:20.0:19.8:20.4. For more on this, see Wooders (2008).

that players are following minimax in roughly 20-45 percent of the cases, compared to 5 percent rejections for the soccer players in Palacios-Huerta and Volij (2007).²¹ Figures 3a and 3b complement these results by providing an ocular depiction of the distributions of individual choice frequencies, compared with the minimax binomial distribution.

On the runs test reported in Panel III of Table 2, our sample of soccer players perform reasonably well, but not quite as well as the players in Palacios-Huerta and Volij (2007). College students do quite poorly, and poker players are in between. In this case, rejections in our data are typically for too many runs (switching too infrequently), consonant with Palacios-Huerta and Volij's rejections.

Overall, as was the case with Hide and Seek, the data suggest that our subjects provide choices that are not as close to minimax predictions as found in Palacios-Huerta and Volij (2007). These results hold whether we examine early or late periods of play, as shown in the Supplementary Appendix. It appears that among our professionals what happens in the field stays in the field, establishing important bounds on the generality of the results of Palacios-Huerta and Volij (2007).

Thus far, we have largely focused on statistical testing, leaving the question of proximity to theory on the sidelines. As can be gleaned from Tables 1 and 2, in aggregate minimax predictions perform relatively well when judged against other social science theories. To provide further descriptive evidence at the individual level, we construct Table 3, which reports on proportions of players who deviated from optimal play. The Table classifies players from each of our subject pools (including bridge players) into "deviations from optimality" bins. For instance, a player who played 26%, 39%, 19%, and

²¹ In the worst cases, individuals are playing Diamond over 60 percent of the time (minimax predicts 40 percent), and in one case a soccer player did not play Heart even a single time.

18% in Four-Card Barry would be classified as a “+/-5%” player due to the 26% being 6% higher than the 20% optimal rate. For comparability purposes, in Panel II of Table 3 we compute similar descriptive statistics using the data from Palacios-Huerta and Volij (2007).

The descriptive statistics provide an indication of the large deviations observed in our data. For example, amongst poker players in the Hide and Seek game, 57% (75%) of the evaders (pursuers) deviate by more than 10% from optimal play. Likewise, for soccer professionals, we observe a substantial fraction of players deviating significantly from optimal play in Four Card Barry: for row (column) players 63% (25%) deviate by more than 10% from optimal play. In contrast, the data in Palacios-Huerta and Volij (2007) are notable in the sense that very few individuals depart significantly from optimal play. For instance, in the Hide and Seek game none of the soccer players deviate by more than 5% from optimal play, and in Four Card Barry only 5% of the soccer professionals deviate by at least 10% from optimal play, and none deviate by more than 15%.

C. Understanding why play deviates from minimax

There are several different plausible explanations as to why subjects in our sample fail to play the minimax strategy. A first explanation is that players would like to play minimax, but they are unable to do so because they cannot solve for the equilibrium, or they are cognitively able but deem the costs too prohibitive. A second, very different explanation, is that they do not believe that their opponents will play minimax. If opponents systematically deviate from minimax (or are expected to deviate), then minimax is no longer the best response because the opponent’s strategy is exploitable. A third explanation lies at the foundation of the experimental environment: the nature and context

of the constructed situation did not induce the professionals to retrieve the relevant cognitive tool kit to play optimally.

We use two approaches to provide a deeper understanding of our results. The first is to use computers as opponents in similar lab games. In one treatment we programmed the computer to play minimax (and then to exploit its competitor if possible), and in the other the computer was programmed to persistently play sub-optimally. Table 3 reports results for the two computer treatments, with the optimally programmed opponent shown in columns 1-2 and the naïve computer opponent in columns 3-4. Results are shown separately for the Hide and Seek and Four-Card Barry games. The top portion of Table 3 presents the same three tests included in the preceding tables, except that we restrict our tests to the behavior of the human player.²²

Importantly for our purposes, even when faced with an opponent programmed to initially play minimax and to only deviate from that strategy in response to non-minimax play by the subject, poker players' actions are not consistent with minimax theory.²³ For a majority of the tests, empirical results are quite similar to the results obtained from the human-human interactions. Indeed, if anything, the runs test reveals that players are more likely to exhibit serially dependent play when competing against the computer, a result consonant with Budescu and Rapoport (1992) if subjects in this treatment interpreted the situation as an individual, non-competitive, choice.²⁴ These results indicate that the

²² As would be expected, the play of the naively programmed computer is nearly always rejected as being consistent with minimax. Less frequently, but often, the play of the optimally programmed computer is also rejected since it deviates from minimax in response to sub-optimal play on the part of the human subject. These results are presented in the Supplementary Appendix.

²³ When we divide each player's actions into two equal size sets corresponding to the first 75 and last 75 plays, we find similar results (see the Supplementary Appendix).

²⁴ In Four-Card Barry, the "world class" poker players perform better than the other players in the sample, although the sample size is small. In Hide and Seek, "world class" players do not play better than the others.

deviations we observe from minimax are not merely due to beliefs that the other player is not playing minimax and can therefore be exploited.

The bottom panel of Table 3 reports other results for these treatments, including average payoffs as well as the fraction of players who “beat” the computer in the sense of winning more than half of the trials.²⁵ As expected, the computer programmed to play optimally slightly outperforms its human opponents: humans win 49.6 (48.5) percent of the payoffs in Hide and Seek (Four-Card Barry). Alternatively, our subjects fared better against the computer programmed to play naïve, non-minimax strategies, particularly in Four-Card Barry, where humans obtained 57.5 percent of the payoffs, and 12 of 13 humans earned more than the computer. This result accords with insights gained in Messick (1967), Fox (1972), Coricelli (2004), and Shachat and Swarthout (2004), who report that subjects have some propensity in these games to exploit non-optimal play.

While our subjects were able to exploit effectively, they did not perform optimally. Given that the naïve computer chose randomly between the four strategies with equal probability, the optimal human strategy is to always play Diamond as the row player and never play Diamond as the column player, yielding an expected payoff of 62.5 percent. None of our subjects realized this payoff level.

In contrast to Four-Card Barry, humans did quite poorly against the naïvely programmed computer in Hide and Seek, winning only 50.9 percent of the total payoff, when the pure-strategy best response to the computer’s strategy would yield an expected payout of 75 percent. Only 8 of 13 humans beat the naïve computer in Hide and Seek.

²⁵ Similar to the human-human treatments, the computer and the player switched roles (row vs. column) halfway through the experiment so that each had the same expected value in terms of wins.

Overall, both sets of results mirror Fox (1972), who finds that subjects adjusted their play in the direction of a best response, but did not play optimally.

Our second approach to learning why minimax met with limited success in our lab experiments is to use post-experimental surveys for the soccer players. We used two survey instruments each given to a part of the professional soccer sample (both are contained in the Supplementary Appendix). One survey asks the question “Does this game remind you of any other games?” None of the twelve soccer players asked this question spontaneously made the link between the experimental game and penalty kicks. Among the remaining soccer players, when prompted for a comparison between the lab game and penalty kicks, 4 players responded that they saw no comparison at all, 2 said that they only thought about penalty kicks *after* the question was asked, 9 said that they were somewhat comparable, and 5 gave an outright yes.²⁶ These results suggest that overall our experiment was not able to summon important field parallels for our subjects.

IV. Conclusions

Determining the conditions under which people play mixed strategies is a question of fundamental importance in economics. Indeed, O’Neill (1991, p. 503) asserts that “the most basic idea in game theory is the minimax argument.” Judged by past laboratory games with unique mixed-strategy equilibria, this most basic idea has been met with limited success. Palacios-Huerta and Volij (2007) have revived this area of research by generating data that is aligned with some of the main predictions of minimax theory.

²⁶ Interestingly, when asked their strategies on penalty kicks, 44 percent of the players reported playing pure strategies (e.g., always kick left), highlighting the fact that very few professional soccer players ever get the chance to take penalty kicks in games and suggesting that few subjects viewed our experimental environment as having a direct parallel to their field of expertise.

In contrast to their results, however, we find that neither professional poker players nor our own sample of professional soccer players actually produce play in the lab that closely approximates minimax predictions. The results of pitting humans against computers programmed to exploit deviations from minimax further demonstrates that the failure to play minimax does not seem to be due to players' beliefs in suboptimal play by their opponents. Although these subjects might be able to randomize effectively in their chosen line of business, they seemingly had difficulty transferring their particular field situation to the specific lab task. Our post-experimental surveys for the soccer treatments highlight the fact that players did not make the connection between the lab experiments and naturally-occurring situations they face in the field.

We should stress that our results are meant to provide important bounds on Palacios-Huerta and Volij (2007), as we note several differences between the two experiments. First, our experiments involved fewer repetitions (75 rounds in each role rather than 200 rounds in a single role), and the subjects were told the number of rounds in advance. Second, we conducted the experiments in the soccer teams' locker rooms, rather than in a university laboratory, and we did not employ screens to hide the backs of a player's cards from his opponent. Third, the experiments were played between teammates, rather than across teams, and we were not able to obtain enough goalkeepers to ensure one goalkeeper per pair. Finally, the subjects were players from American professional teams rather than from Spanish professional teams. We do not know which, if any, of these differences might have caused the large behavioral discrepancies, but one point of the study

is that the previous results on soccer players are not as robust as one might have hoped. Additional replications would be valuable.²⁷

Clearly, subjects come to experiments with rules of behavior learned in the outside world. Depending on whether the specific context of the lab game cues the proper rules of thumb, radically divergent results can be obtained. Harrison and List (2007), for instance, examine the behavior of professional bidders in their naturally occurring environments. In their real-world bidding, such subjects do not constantly fall prey to the winner's curse. When the expert bidders are placed in unfamiliar roles, however, they often fall prey to the winner's curse, just as happens in the lab. Our results combined with their insights underscore an important methodological point: slight changes in context can have profound behavioral effects, whether students or professionals are the experimental participants.

²⁷ We attempted to reproduce our results with American college soccer players, but quickly realized that this would violate NCAA rules prohibiting payments to college athletes. One coach suggested that players could legitimately compete if the cash they earned would be donated to charity, but we felt this would not be a very good test of equilibrium behavior in a zero-sum game.

References

- Brayer, A.R. 1964. An experimental analysis of some variables of minimax theory. *Behavioral Science* 9 (1): 33–44.
- Brown, J.N., and R.W. Rosenthal. 1990. Testing the minimax hypothesis: A re-examination of O'Neill's game experiment. *Econometrica* 58(5): 1065–1081.
- Budescu, D.V., and A. Rapoport. 1992. Generation of random series in two-person strictly competitive games. *Journal of Experimental Psychology* 121 (3): 352-363.
- Chiappori, P.A., S. Levitt, and T. Groseclose. 2002. Testing mixed strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer. *American Economic Review* 92 (4):1138–1151.
- Coricelli, G. 2004. Strategic interaction in iterated zero-sum games. Unpublished Manuscript. http://economics.eller.arizona.edu/downloads/working_papers/coricelli.pdf
- Fehr, E., and J.A. List. 2004. The hidden costs and returns of incentives: Trust and trustworthiness among CEOs. *Journal of the European Economic Association* 2 (5): 743-71.
- Fox, J. 1972. The learning of strategies in a simple, two-person zero-sum game without saddlepoint. *Behavioural Science*. 17 (3): 300–308.
- Friedman, L. 1971. Optimal bluffing strategies in poker. *Management Science*, 17 (12): B764-B771.
- Gibbons, J.D., and S. Chakraborti. 2003. *Nonparametric statistical inference, 4th edition*. New York: Marcel Dekker, Inc.
- Haigh, M., and J.A. List. 2005. Do professional traders exhibit myopic loss aversion? An experimental analysis. *Journal of Finance* 60 (1): 523-534.
- Harrison, G.W. and J.A. List, 2004. Field experiments. *Journal of Economic Literature* 42, 1009-1055.
- Harrison, G.W., and J.A. List. 2007. Naturally occurring markets and exogenous laboratory experiments: A case study of the winner's curse. *Economic Journal* (forthcoming)
- Hirschberg, Daniel, Steven Levitt, and John List. 2008. Poker Players Equate Mixed-Strategy Payoffs in the Field. Unpublished manuscript.
- Hsu, S-H., C-Y. Huang, and C-T. Tang. 2007. Minimax play at Wimbledon: Comment. *The American Economic Review* 97 (1): 517-523.

- Lieberman, B. 1960. Human behavior in a strictly determined 3×3 matrix game. *Behavioral Science* 5(4): 317–322.
- Lieberman, B. 1962. “Experimental studies of conflict in some two-person and three-person games.” In *Mathematical Models in Small Group Processes*, edited by J.H. Criswell, H. Solomon, and P. Suppes. Stanford, CA: Stanford University Press. 203–220.
- Loewenstein, G. 1999. Experimental economics from the vantage-point of behavioral economics.” *Economic Journal* 109(453): F23-34
- Messick, D.M. 1967. Interdependent decision strategies in zero-sum games: A computer-controlled study. *Behavioral Science* 12 (1): 33–48.
- O'Neill, B. 1987. A non-metric test of the minimax theory of two-person zerosum games. *Proceedings of the National Academy of Sciences* 84 (7): 2106-2109.
- O'Neill, B. 1991. Comments on Brown and Rosenthal’s reexamination. *Econometrica* 59 (2): 503-507.
- Palacios-Huerta, I. 2003. Professionals play minimax. *Review of Economic Studies* 70 (2): 395-415.
- Palacios-Huerta, I. and O. Volij. 2008. Experientia docet: Professionals play minimax in laboratory experiments. *Econometrica* 76: 71-115.
- Rosenthal, R.W., J. Shachat, and M. Walker. 2003. Hide and seek in Arizona. *International Journal of Game Theory* 32 (2): 273-293.
- Shachat, J., and T.J. Swarthout. 2004. Do we detect and exploit mixed strategy play by opponents? *Mathematical Methods of Operations Research* 59 (3): 359-373.
- Spiliopoulos, L. 2007. Do repeated game players detect patterns in opponents? Revisiting the Nyarko and Schotter belief elicitation experiment. *Munich Personal RePEc Archive* Paper No. 2179. <http://mpira.ub.uni-muenchen.de/2179>
- Von Neumann, J. 1928. Zür theorie der Gesellschaftsspiele. *Mathematische Annalen* 100 (1): 295–320.
- Walker, M., and J. Wooders. 2001. Minimax play at Wimbledon. *American Economic Review* 91 (5): 1521–38.
- Wagenaar, W.A. 1971. Serial non-randomness as a function of duration and monotony of a randomization task. *Acta Psychologica* 35 (1): 78–87.
- Wooders, J. 2008. Does Experience Teach? Professionals and Minimax Play in the Lab, unpublished manuscript.