Learning to Make Better Strategic Decisions*

Eric Cardella†

June 10, 2011

Abstract

Strategic settings are often complex and agents who lack deep reasoning ability may initially fail to make optimal decisions. This paper experimentally investigates how the decision making quality of an agent's opponent impacts learning-by-doing (LBD) and learning-by-observing (LBO) in a strategic decision making environment. Specifically, does LBD become more effective when agents face opponents who exhibit optimal decision making? Similarly, does LBO become more effective when agents observe opponents who exhibit optimal decision making? I consider an experimental design that enables me to measure strategic decision making quality and control the decision making quality of an agent's opponent. The results suggest that LBD is more effective when facing an optimal decision making opponent. However, LBO does not appear to be more effective when observing an optimal decision making opponent. The results also suggest that LBD and LBO are equally effective in improving decision making quality in the game considered, irrespective of the opponent's decision making quality.

*I would like to thank my dissertation advisor Martin Dufwenberg for helpful comments. I would also like to thank Anna Breman, David Butler, Tamar Kugler, John Wooders, and conference participants at the University of Arizona for additional helpful comments. I am greatful to the Economic Science Laboratory at the University of Arizona for providing financial support.

†Department of Economics, University of Arizona, McClelland Hall 401, PO Box 210108, Tucson, AZ 85721-0108; Email: ecardell@email.arizona.edu.

1 Introduction

Economic settings are often complex, and optimal decision making can require deep reasoning ability. Oligopolies, negotiations, contracting, and auctions represent a few of the many complex economic settings where agents are called upon to make important decisions. Agents who lack high levels of strategic sophistication and/or deep reasoning ability are likely to make sub-optimal decisions, which can often lead to inefficient outcomes. As a result, investigating how agents learn to make better decisions remains an important and largely open research question. The goal of this paper is to provide insights regarding how agents become better strategic decision makers.

One learning mechanism that can facilitate improved decision making is learning-by-doing (LBD). Through repeated experience, agents acquire knowledge and skill that can subsequently lead to better decision making. For example, in a seminal paper, Arrow (1962) argues that "learning is the product of experience. Learning can only take place through the attempt to solve a problem and therefore only takes place during activity" (p. 155). The body of literature related to LBD is extensive, and I refer the reader to Thompson (2010) for a comprehensive review of theoretical applications and empirical investigations of LBD. An alternative, yet related, learning mechanism that can facilitate improved decision making is learning-by-observing (LBO). By observing the decision making of others, agents acquire knowledge and skill that can subsequently lead to better decision making. For example, Jovanovic and Nyarko (1995) present a model of LBO where an "apprentice" learns from the skillful "foreman" he observes. Merlo and Schotter (2003) and Nadler et al. (2003) provide experimental evidence that supports LBO.[1]

The previous literature related to LBD and LBO has focused primarily on learning in single agent decision tasks. However, strategic settings are often characterized by decision making of multiple agents. In these strategic settings, the process of acquiring knowledge and skill through LDB or LBO for an agent will likely be influenced by the decision making of the other agents. The primary motivation of this paper is to experimentally investigate how LBD and LBO are influenced by the nature of the decision making quality of an agent's opponent. Specifically, I investigate whether LBD becomes more effective when an agent repeatedly play against an opponent who makes optimal decisions, compared to sub-optimal decisions. Similarly, I investigate whether LBO becomes more effective when an agent repeatedly observe an opponent who makes optimal decisions, compared to sub-optimal decisions.

To shed light on these questions, I propose a stylized experimental design that enables me to identify and measure decision making quality in a strategic setting. Subjects repeatedly play a 2-player, sequential-move game, which features a dom-

[1] LBO has also been well documented in several animal experiments including John et al. (1969), Tomasello et al. (1987), and Terkel (1996).

inant strategy. A full description of the game follows in the subsequent section. The dominant strategy of the chosen game serves as an identifiable and measurable proxy for optimal decision making. The design also features the implementation of pre-programmed computer opponents, which enables me to explicitly control the decision making quality of each subject's opponent.[2] Specifically, I consider two types of computer opponents: The first, which I refer to as the *optimizing opponent,* is pre-programmed to play a dominant strategy, i.e., make optimal decisions. The second, which I refer to as the *naïve opponent,* is pre-programmed not to play a dominant strategy, i.e., make sub-optimal decisions. This variation in the opponents decision making quality is what allows me to identify how LBD and LBO in strategic settings are impacted by the nature of the opponent's decision making quality.

As a second motivation, the experimental design allows me to experimentally compare the effectiveness of LBD and LBO in a strategic setting. This is similar in spirit to Merlo and Schotter (2003) who experimentally compare LBD and LBO in a single-agent profit maximization problem.[3] In this single agent decision task, the authors find evidence that LBO is more effective than LBD. This paper differs from Merlo and Schotter in two important ways: First, I consider the comparison of LBD with LBO in a 2-player dynamic game. Second, I compare LBD with LBO for both the *optimizing* and *naïve opponent*. The comparisons between LBD and LBO in this study can be viewed as complementary to that of Merlo and Schotter (2003).

The motivation of this study is to gain valuable insights regarding how agents learn to make better decisions in strategic settings. Ideally, one would want to gain such insights by experimenting in real strategic settings in the field. However, identifying and measuring decision making quality in the field is difficult. Optimal decisions often depend on the agent's preferences and/or beliefs, both of which are difficult to measure. Without information regarding an agent's preferences and/or beliefs, it may be difficult to measure the quality of an agent's decisions in the field. A similar line of reasoning applies to the difficulty of identifying the decision making quality of an agent's opponent in field settings. However, by considering an experimental game that features a dominant strategy, I am able to identify and measure optimal strategic decision making. Additionally, the lab provides a setting where it is possible to systematically control the decision quality of an agent's opponent by using pre-programmed computer opponents. Therefore, a stylized lab experiment enables me

[2]The use of pre-programmed computer opponents is certainly not novel to this study. Previous studies include Johnson et al. (2002) who use pre-programmed computer opponents in an alternating bargaining game. Merlo and Schotter (2003) use pre-programmed computer opponents in a tournament game. Shachat and Swarthout (2004) use pre-programmed computer opponents in a matching pennies game. Dursch et al.(2010) use pre-programmed computer opponents in a repeated Cournot game.

[3]Technically, the authors consider a 2-player simultaneous move tournament game. However, the authors effectively transform the 2-player game into a single agent profit maximization problem by informing subjects that they will face a computer opponent that is pre-programmed to make a certain pre-specified decision.

to gain valuable insights regarding the impact of the opponents decision making quality on LBD and LBO in strategic settings, which would otherwise be difficult by experimenting in the field.

I find experimental evidence that LBD is significantly more effective when subjects play the *optimizing opponent*, relative to the *naïve opponent.* However, there is little evidence that LBO is significantly more effective when subjects observe the *optimizing opponent,* relative to the *naïve opponent.* I also find very little difference in decision making quality between subjects who initially play and subjects who initially observe, irrespective of the decision making quality of the opponent. The results of this paper suggest that LBD and LBO can be comparable mechanisms for becoming a better strategic decision maker. Additionally, the effectiveness of LBD in strategic settings can be increased by playing an optimal decision making opponent, whereas the effectiveness of LBO seems unaffected by the decision making quality of the opponent. Although the experimental design features a stylized strategic settings, the results from this study can be useful for informing a broader range of topics including, but not limited to: the design of effective teaching modules, the design of effective job training programs, or perhaps on a more casual level, a lucrative poker career.

The fact that the chosen experimental game features a dominant strategy makes it atypical of many strategic settings in the field. This may be particularly salient in light of the results from several studies that have shown differences in learning between intellective tasks and judgement tasks (Hill (1982), Hastie (1986), Levine and Moreland (1998), and Laughlin et al. (2002) for a thorough review). Intellective tasks are characterized as featuring a demonstrably correct solution, whereas judgement tasks are evaluative and do not feature a demonstrably correct solution (Laughlin (1980), Laughlin and Adamopoulos (1980), and Laughlin and Ellis (1986)). Games that feature a dominant strategy, like the one considered in this study, are similar to intellective task in the sense that there exists a demonstrably correct solution. I acknowledge that this may limit the generalizability (Levitt and List (2007)) of the insights gleaned from the experimental results regarding learning in general strategic settings. Nevertheless, even at its most limited scope, this study can help better our understanding of LBD and LBO in strategic settings that feature a dominant strategy. Such insights can be valuable when applied to games in the lab featuring dominant strategies, e.g., 2-player guessing games (Grosskopf and Nagel (2008)), and games in the field featuring dominant strategies, e.g., second-price auctions.[4]

The paper proceeds by formally describing the experimental design and developing the research hypotheses in Section 2. The results are presented in Section 3, and I conclude with discussion in Section 4.

[4] Grosskopf and Nagel (2008) speak to the importance of investigating learning in games with dominant strategies by noting in their conclusion that "it remains to be investigated whether and how people can learn to choose zero in the n=2 BCG" (p. 98).

2 Experimental Design and Research Hypotheses

2.1 The Game of 21

I consider a 2-player, sequential move, constant-sum game of perfect information. The game, introduced for experimental purposes by Dufwenberg et al. (2010) (DSB henceforth), is the game of 21 (G21 henceforth). This game begins with the first mover choosing either 1 or 2. After observing the first mover's choice, the second mover chooses to increase the count by either 1 or 2. For example, if the first mover chooses 2, the second mover can choose either 3 or 4. The players continue alternating turns increasing the count by either 1 or 2; the player who chooses 21 wins.[5] What is the optimal way to play G21?

Upon some reflection, one realizes that G21 features a second mover advantage, where the second mover can guarantee victory by choosing every multiple of three, i.e., choosing 3, 6, 9, 12, 15, 18, and then 21 to win the game. Thus, any strategy that prescribes choosing every multiple of three is dominant for the second mover.[6] In addition, any subgame that contains an available multiple of three also features a dominant strategy of choosing that multiple of three followed by all subsequent multiples of three. That is, if one of the players fails to choose a multiple of three, the other player could guarantee victory by choosing that multiple of three and all subsequent multiples of three. I generally refer to the class of strategies that prescribes choosing every available multiple of three, which guarantees victory, as the dominant solution to G21.

G21 features two properties that make it well-suited for the purposes of this study. First, G21 features a dominant solution that acts as an identifiable and measurable proxy for optimal decision making in G21. Simultaneously, the dominant solution eliminates the ambiguities in identifying optimal decision making that can arise from differences in beliefs about the opponent's strategy. Second, G21 is constant-sum and

[5] Gneezy et al. (2010) concurrently introduced two related games for experimental purposes, which the authors refer to generally as "race games". In their G(15,3) (G(17,4)) race game, subjects alternate incrementing the count by 1, 2, or 3 (1, 2, 3, or 4), and the first person to reach 15 (17) wins. Levitt et al. (forthcoming) consider two versions of the "race to 100" game where subjects alternate incrementing the count by either 1-9 or 1-10 and the person to reach 100 wins. Dixit (2005) refers to an empirical account of a related game, the 21 flags game, which appeared on an episode of the TV series "Survivor" in 2002 as an immunity challenge between two teams. The two teams alternated removing 1, 2, or 3 flags from an initial pile of 21 flags. These related "race" games feature a similar structure and strategic properties to those of G21.

[6] Choosing every multiple of three does not describe a complete strategy, as it does not specify an action for the agent at information sets where a multiple of three is not available. A complete strategy must specify an action for the second mover at these information sets, although these information sets will not be reached if the second mover plays a dominant strategy of choosing every multiple of three. Any strategy that specifies choosing every available multiple of three, regardless of what is chosen when a multiple of three is not available, is dominant because it will guarantee victory for the second mover. Thus, the second mover has many dominant strategies.

features a binary outcome of "win" or "lose". This eliminates possible ambiguities in identifying optimal decisions that can arise from efficiency concerns, distributional preferences, and/or belief based motivations. Differences in beliefs about others, and the presence of social preferences can lead to obvious confounds when trying to identify optimal decision making in other frequently implemented experimental games, e.g., three or more players guessing (p-beauty contest) games, trust games, centipede games, and sequential bargaining games.[7]

In order to investigate LBD and LBO, it is crucial that the chosen game be complex enough to allow for learning. That is, G21 must be complex enough that most subjects playing G21 for the first time fail to play the dominant solution. DSB (2010) find that roughly 85 percent of subjects playing G21 for the first time initially fail to play the dominant solution. Similarly, Gneezy, Rustichini, and Vostroknutov (2010) also find that most subjects initially fail to play the dominant solution in the related race games they consider. However, in both studies, many subjects learn to play the dominant solution after several rounds of play. The previous experimental results suggest that G21 is sufficiently complicated that most subjects initially fail to play a dominant solution, yet it is sufficiently straightforward to admit the possibility of improved decision making. Additionally, the sequential structure of G21, compared to a simultaneous move game, allows for the construction of more robust measures of decision making quality. These measures enable me to better quantify the degree to which subjects are deviating from optimal decision making.

G21 is a stylized strategic game, and this feature creates a clear cost of the experimental design. Namely, the fact that G21 features a dominant solution may limit the extent to which the experimental insights apply to a general class of strategic settings. However, the stylized features of G21 create a clear benefit of the experimental design. Namely, the fact that G21 features a dominant solution, is constant sum, and features a binary outcome enables me to identify, and measure, strategic decision making quality. Thus, G21 provides a suitable platform for gaining initial insights regarding how the decision making quality of an opponent affects LBD and LBO in strategic settings.[8]

[7]See Nagel (1995), Duffy and Nagel (1997), Bosch-Domenech et al. (2002), and Ho et al. (1998) for experimental studies using the guessing game. See Rosenthal (1981), McKelvey and Palfrey (1992), Fey et al. (1996), Rapoport et al. (2003), and Palacios-Huerta and Volij (2009) for experimental studies using the centipede game. See Binmore et al. (1985), Ochs and Roth (1989), Harrison and McCabe (1992, 1996), Johnson et al. (2002), Binmore et al. (2002), and Carpenter (2003) for experimental studies of sequential bargaining.

[8]The use G21, in lieu of other variations of the general class of race games, was strictly a design choice. All variations of the race game feature similar properties: a dominant solution, constant sum, and a binary outcome. Therefore, any other variation of the race game would have been equally suitable.

2.2 Experimental Treatments and Procedure

The subject pool for this experiment consisted of undergraduates from the University of Arizona. All of the sessions were conducted in the Economic Science Laboratory (ESL) at the University of Arizona. A total of four sessions were conducted using 96 subjects. This was a computerized experiment, and the experimental software was programmed in Microsoft Visual Basic. A copy of the experimental instructions and sample screen shots are presented in the Appendix.

When playing G21, subjects were always matched with a pre-programmed computer opponent. There were two possible types of computer opponents, which I refer to as the *optimizing opponent* and the *naïve opponent*. The *optimizing opponent* was pre-programmed to play the dominant strategy of choosing every available multiple of three, and randomly increment the count by 1 or 2 when a multiple of three was not available. The *naïve opponent* was pre-programmed to randomly increment the count by 1 or 2 at every decision. All subjects were informed that they were playing against a computer opponent, but unaware of which type of opponent.[9]

I implement a between-groups design with four treatments. Each of the four treatments are described as follows:

Optimal Player The subject played six rounds of G21 against the *optimizing opponent* followed by six rounds of G21 against the *naïve opponent*.

Naïve Player The subject played all twelve rounds of G21 against the *naïve opponent*.

Optimal Observer The subject first observed a subject from the Optimal Player treatment play the initial six rounds of G21, and then proceeded to play six rounds of G21 against the *naïve opponent*.

Naïve Observer The subject first observed a subject from the Naïve Player treatment play the initial six rounds of G21, and then proceeded to play six rounds of G21 against the *naïve opponent*.

Twenty four subjects participated in each of the four treatments. Each experimental session consisted of an equal number of randomly assigned players and observers. Each player was instructed to sit at their assigned computer carrel and each observer was instructed to stand quietly behind their assigned player. An experimenter was

[9] Hall (2009) implements a similar design where subjects play variations of the race game against a pre-programmed computer opponent in order to experimentally investigate how monetary incentives affect a subject's ability to backward induct. The author considers two variations. In the first, subjects race to 21 incrementing the count by 1-3, and in the second, subjects race to 21 incrementing the count by 1-4. However, Hall only considers a pre-programmed opponent who always plays the dominant strategy when possible and randomizes otherwise, which is analogous to what I refer to as the *optimal opponent*.

present to monitor all sessions and to ensure that there was no communication between players and observers. After reading the instructions, players proceeded to play six rounds of G21 against their assigned computer opponent, and observers merely observed the play.[10] After observing the first 6 rounds of play, the observers were then instructed to move to their own carrel where they proceeded to play six rounds of G21. The players then proceeded to play an additional six rounds of G21 with no observer. All subjects were allowed to progress at their own pace to ensured that each subject had ample time to make well thought out decisions without time pressure. Each subject began as the first mover and alternated between first mover and second mover in all subsequent rounds.

Before I proceed, let me first highlight the motivation to have all four treatments play the final six rounds of G21 against the *naïve opponent*. First, the treatments only differ in terms of the subjects experience during the first six rounds of play. As a result, any differences in decision making between the treatments in the last six rounds can then be attributed to the subjects' experience in the first six rounds. This isolates the effect of the decision making quality of the opponent on LBD and LBO, and thus answers the research questions of this study.

Second, playing the final six rounds against the naïve opponent allows for the construction of a variety measures (described in the next section) that can be used to compare decision making quality between treatments. In rounds where subjects play the *optimizing opponent* very little insight can be gained regarding a subject's decision making quality. When subjects act as the first mover against the *optimizing opponent*, they will not have an opportunity to play the dominant solution. Similarly, when subjects act as the second mover and fail to choose a multiple of three, they will be unable to begin playing the dominant strategy in any later subgame. However, when subjects play against the *naïve opponent*, opportunities will likely arise at various subgames to begin playing a dominant strategy.[11] As a result, requiring all subjects to play against the *naïve opponent* in the last six rounds allows me to investigate multiple dynamic measures of decision making quality in G21. This enables me to better quantify the extent that subjects are deviating from optimal decision making, which provides a more robust conclusion about the impact of facing an opponent who makes optimal decisions on the effectiveness of LBD and LBO.

[10] It is possible that having an observer can affect the decision making of the player. However, an observer was present in both the optimal and naïve treatment which controls for this possible effect, if it exists. Additionally, there is little reason to think that if there is an observer effect that it is systematically different between the treatments; therefore, any observer effect is not likely to bias the results when investigating the relative difference between the optimal player and naïve player treatments.

[11] Because the *naive opponent* randomizes equally at every decision node, if a subject fails to choose an available multiple of three, then there is a 50% chance that a multiple of three will be available on the subjects next turn.

All subjects were paid a $5 USD show up fee. Additionally, the players received $1 USD for every round they won for all twelve rounds. The observers were paid $1 USD for every round that their corresponding player won during the first six rounds plus $1 USD for every round they won while playing during the last six rounds. Paying the observer for the first six rounds while they observed mitigated any spiteful feelings that may arise from possible inequity in payments between player and observer. In addition, this controlled for the amount of money won in the first 6 rounds between the player and observer. This ensured that any differences in behavior between the two treatments was not systematically caused by differences in monetary earnings during the first 6 rounds. Each session lasted about thirty minutes and no subjects participated in more than one session.

2.3 Research Hypotheses

The two primary research question of this study are whether (1) LBD becomes more effective when an agent repeatedly plays against an optimal decision making opponent, and whether (2) LBO becomes more effective when an agent repeatedly observes an optimal decision making opponent, compared to sub-optimal decision making opponents. I attempt to answer these questions using the experimental data to test the following two corresponding hypotheses:

H1: Subjects from the Optimal Player treatment exhibit better decision making than subjects from the Naïve Player treatment.

H2: Subjects from the Optimal Observer treatment exhibit better decision making than subjects from the Naïve Observer treatment.

The secondary objective of this study is to compare LBD with LBO when a subject faces an *optimizing opponent* and a *naïve opponent*. I test these comparisons between LBD and LBO with the following two corresponding hypotheses:

H3: Subjects from the Optimal Player treatment exhibit equivalent decision making quality as subjects from the Optimal Observer treatment

H4: Subjects from the Naïve Player treatment exhibit equivalent decision making quality as subjects from the Naïve Observer treatment

Testing H1-H4 requires measuring the decision making quality of subjects playing G21. To help minimize measurement error and provide a more robust test of the research hypotheses, I consider several complementary measures that characterize decision making quality in G21. Additionally, I focus on aggregate comparisons between treatments to help smooth the noise resulting from subject heterogeneity in regards to inherent decision making aptitude. When testing the research hypotheses, I only compare decision making for the last 6 rounds of play, identified in the data as rounds 7-12, when subjects played G21 against the *naïve opponent*. Therefore, any differences in decision making quality between treatments can be attributed to the subject's treatment during the initial 6 rounds, namely whether the subject initially played or observed, and whether the opponent was optimal or naïve.

3 Results

Testing H1-H4 is only possible conditional on the maintained assumption that G21 is sufficiently complex that most subjects initially fail to make optimal decisions. Hence, the possibility for learning, and the consequent ability to test how LBD and LBO are affected by the decision making quality of a subject's opponent. To verify this assumption, I look at the data from round 2; the first round when subjects acted as the second mover and, therefore, had an opportunity to play the dominant solution. The data confirms that most subjects initially fail to make optimal decisions in G21. In particular, only 5/24 (21%) subjects from Optimal Player treatment, and 2/24 (8%) subjects from the Naïve Player treatment played the dominant solution in round 2. Given that most subjects initially fail to play the dominant solution, it is possible to proceed with testing H1-H4.

3.1 Testing H1: Optimal Player Treatment vs Naïve Player Treatment

In this section, I compare the decision making patterns from the Optimal Player and Naïve Player treatments to test H1, namely subjects from the Optimal Player treatment exhibit better decision making than subjects from the Naïve Player treatment. The first measure of decision making quality I consider is a *mistake free game* (MFG). I define a MFG as a game of G21 where the subject chooses every available multiple of three. In rounds where the subject acts as the second mover, this corresponds to choosing every multiple of three.[12] In rounds where the subject acts as the first mover, this corresponds to choosing the first available multiple of three and every subsequent multiple of three. Because playing a MFG requires choosing several sequential multiples of three, it is unlikely that a subject would just randomly play a MFG.[13] Thus, playing more MFGs can be viewed as evidence consistent with better decision making in G21. Table 1 reports the aggregate proportion of MFGs for each

[12] When subjects act as the second mover, the measure of MFG is analogous to the measure of "perfect play" considered by DSB (2010).

[13] In order for a subject to play a MFG when they act as the second mover, they must choose 7 sequential multiples of three. The probability of playing a MFG if the subject was randomly incrementing the count by one or two is $\frac{1}{2^7} \approx .008$. In addition, there is no reason to suspect that the likelihood of randomly playing a MFG is different in the two treatments.

round and the average total number of MFGs, aggregated over the last six rounds, per player.

Table 1
MFG Comparison between Optimal Player and Naïve Player

	Aggregate Proportion of MFG	
Round	Optimal Player	Naïve Player
7	4/24	6/24
8	9/24**	3/24
9	6/24	5/24
10	8/24*	3/24
11	8/24	4/24
12	8/24	6/24
Total MFG / Player	1.79/6	1.13/6

Notes: The proportions of MFGs were tested using 1-sided Fisher's Exact test.
The Total MFG/Player was tested using a 1-sided Mann-Whitney test.
* significantly greater at the 10% level ** significantly greater at the 5% level

Table 1 reveals that in rounds 8-12 the proportion of MFGs is higher in the Optimal Player treatment. For each round, a 1-sided Fisher's Exact test was used to test the null hypothesis that the proportion of MFGs was equal for the two treatments, against the alternative that the proportion of MFGs is greater in the Optimal Player treatment. (see Siegel and Castellan (1988)). From Table 1, we can see that we reject the null in round 8 ($p = 0.047$) and round 10 ($p = 0.084$).[14] Aggregated over the last 6 rounds, subjects from the Optimal Player treatment played, on average, 1.79/6 MFGs compared to 1.13/6 in the Naïve Player treatment. However, this difference is not statistically significant using a 1-sided Mann-Whitney U-test. The data reveals that subjects from the Optimal Player treatment play marginally more MFGs than subjects from the Naïve Player treatment; which provides some initial evidence consistent with H1.

The data from table 1 reveals that only a relatively small fraction of the total games were played mistake free. This highlights the importance of investigating how agents become better decision makers when they might lack the strategic sophistication to make optimal decisions. Subjects may fail to play a MFG, but nevertheless,

[14] All proportions tests for this study were done with a Fisher's Exact test. However, the significance levels for all proportions tested are robust to using a Pearson Chi-Squared test.

exhibit some improved decision making. For example, one could argue that a subject who failed to choose 3 (when it was available), but proceeded to choose every subsequent available multiple of three in the game, exhibited better decision making than a subject who failed to choose all available multiples of three in a game. To get a more complete comparison of decision making quality between the two treatments, I consider two additional dynamic measures of decision making quality.

The second measure I consider is error of type-n (E-n). I define E-n as a failure to choose the available nth-multiple of three in a given game of G21.[15] For example, a subject makes an E-6 if he failed to choose 6, when it was available. Given that choosing every available multiple of three guarantees victory, failing to choose an available multiple of three can be viewed as a sub-optimal decision in G21. The third measure I consider is last error (LE). I define LE as the last multiple of three, in a given game of G21, that a subject explicitly failed to choose. For example, if a subject failed to choose 6, when it was available, but proceeded to choose every subsequent available multiple of three for that game, then for that game the subject would record a LE = 6 for that game of G21. The LE measure is an element of the set {0, 3, 6, 9, 12, 15, 18}.[16] The LE measure can be viewed as a proxy for the reasoning ability exhibited by a particular subject in a particular game of G21.[17] Thus, a lower LE level indicates better decision making G21. Table 2 reports the mean per player E-n rates and the Any Error rate, aggregated over the last six rounds rounds. The Any Error measure is defined as the failure to choose *any* available multiple of three.

Table 2 reveals that the mean E-n is lower in the Optimal Player treatment than the Naïve Player treatment for every E-n considered, except E-18. For each E-n, a 1-sided Mann-Whitney U-test was used to test the null hypothesis that the mean E-n rates are equal, against the alternative that the mean E-n rate is lower in the Optimal Player treatment. The null can be rejected for E-3 (p = 0.017), E-9 (p = 0.064), and E-12 (p = 0.060). Additionally, the mean Any Error rate is 29% for subjects in the Optimal Player treatment compared to 36% for the Naïve Player treatment. The null hypothesis that the mean Any Error rates are equal can be rejected in favor of the alternative that the mean Any Error rate is lower for the Optimal Player treatment (p = 0.052). The data reveals that subjects in the Optimal Player treatment make significantly fewer errors in G21, which provides further evidence consistent with H1.

[15] This measure is analogous to the "error of type-n" considered by Gneezy et al. (2010).

[16] It is possible for a subject to fail to choose 21 when it is available. However, this only happened 8 times in 576 games. It is likely that these few instances were accidental and resulted from the specific software design of the user interface (see the Appendix for a sample screen shot). As a result, I drop all 8 cases from the analysis where LE = 21.

[17] I stress the importance that LE is a proxy for number of steps of reasoning. Consider a subject who fails to choose 9 and then just randomly choose 12, 15, 18, and 21. The subject may not have exhibited any steps of reasoning, but nevertheless, LE = 12 is recorded for this subject. Or alternatively, suppose the first available multiple of three comes at 15. If the subject chooses 15, 18, and 21, then LE = 0 for that game which could bias the actual measure of LE downward. However, I consider differences in LE between treatments to smooth such measurement biases.

TABLE 2
Error Rate Comparison Between Optimal Player and Naïve Player

	Per Player Mean Error Rate	
Measure	Optimal Player	Naïve Player
E-3	26%**	47%
E-6	41%	43%
E-9	28%*	40%
E-12	36%*	50%
E-15	23%	33%
E-18	16%	11%
Any Error	29%*	36%

Notes: All error rate measures were tested using a 1-sided Mann-Whitney test
* significantly smaller at the 10% level ** significantly smaller at the 5% level

Table 3 reports the mean LE for rounds 8, 10, 12, and the mean LE per subject aggregated over those rounds. I only consider rounds 8, 10, and 12 because these are the rounds where the subject acted as the second mover. I do this to eliminate possible downward bias in LE that can result when subjects act as the first mover.[18] From Table 3 we can see that the mean LE is lower in the Optimal Player treatment in each of the rounds considered. The null hypothesis that the mean LE is equal between the two treatments can be rejected in round 8 ($p = 0.099$) and round 10 ($p = .043$) using a 1-sided Mann-Whitney U-test. The null hypothesis that the mean aggregate LE of 6.71 in the Optimal Player treatment is equal to the mean aggregate LE of 8.71 in the Naïve Player treatment can also be rejected using a 1-sided Mann-Whitney U-test ($p = 0.084$).

The experimental data is generally consistent with H1. Subjects in the Optimal Player treatment appear to exhibit better decision making in G21 than subjects in the Naïve Player treatment. Specifically, subjects in the Optimal Player treatment tend to play more MFGs, make fewer errors, and make their last error at an earlier subgame than subjects in the Naïve Player treatment. For most measures considered, the difference between the Optimal and Naïve Player treatments was in the direction consistent with H1 and many were significant at standard levels. Taken together, the experimental results suggest that LBD can be more effective in strategic settings when

[18]For example, when a subject acts as the first mover, the computer could proceed to select 3, 6, 9, 12, and then 14. A subject could then select 15, 18, and 21 to win. In this case, the subject chose every available multiple of three and would record a LE = 0.

agents play an optimal decision making opponent, relative to a sub-optimal decision making opponent.

Table 3
LE Comparison Between Optimal Player and Naïve Player

	Per Player Mean LE	
Round	Optimal Player	Naïve Player
8	7.38*	9.38
10	6.50**	9.50
12	6.25	7.25
Aggregate LE	6.71*	8.71

Notes: All LE measures were tested using a 1-sided Mann-Whitney test
* significantly smaller at the 10% level ** significantly smaller at the 5% level

3.2 Testing H2: Optimal Observer Treatment vs Naïve Observer Treatment

In this section, I compare the decision making from the Optimal Observer and Naïve Observer treatments to test H2, namely subjects from the Optimal Observer treatment exhibit better decision making than subjects from the Naïve Observer treatment. I will use the same decision making quality measures described from the previous section. Table 4 reports the aggregate proportion of MFGs for each round, and the average total number of MFGs per player.

From Table 4, we see that the proportion of MFGs is higher in every round for the Optimal Observer treatment compared to the Naïve Observer treatment. We can reject the null hypothesis that the proportion of MFGs is equal for the two treatments in round 8 ($p = 0.055$) and round 12 ($p = 0.068$) using a 1-sided Fisher's Exact test. Additionally, subjects from the Optimal Observer treatment play, on average, 1.79/6 MFGs compared to .83/6 for the Naïve Observer treatment; we can reject the null hypothesis that these averages are equal using a 1-sided Mann-Whitney U-test ($p = 0.050$). The data suggests that subjects from the Optimal Observer treatment play more MFGs than subjects from the Naïve Observer treatment.

Table 4
MFG Comparison between Optimal Observer and Naïve Observer

	Aggregate Proportion of MFG	
Round	Optimal Observer	Naïve Observer
7	8/24	5/24
8	4/24*	0/24
9	10/24	5/24
10	6/24	4/24
11	8/24	4/24
12	7/24*	2/24
Total MFG / Player	1.79/6**	.83/6

Notes: The proportions of MFGs were tested using 1-sided Fisher's Exact test.
The Total MFG/Player was tested using a 1-sided Mann-Whitney test.
* significantly greater at the 10% level ** significantly greater at the 5% level

Table 5 reports the mean per player error rates, aggregated over all rounds, for each treatment. We can reject the null hypothesis that the mean error rates are equal across the two treatments for E-3 (p = 0.046) and E-12 (p = 0.027) using a 1-sided Mann-Whitney U-test. However, we fail to reject the null hypothesis that the mean Any Error rate is equal between the Optimal and Naïve Observer treatments. In terms of comparing LE, Table 6 reveals that for all rounds, we can not rail to reject the null hypothesis that the mean LE is equal between the two treatments. Similarly, we fail to reject the null hypothesis that the aggregate mean LE of 9.67 for the Optimal Observer treatment is equal to the 10.08 aggregate mean LE for the Naïve Observer treatment.

Overall, the experimental data comparing the decision making quality between the Optimal Observer treatment and Naïve Observer treatment yields mixed results. The data reveals that subjects from the Optimal Observer treatment play more MFGs and make marginally fewer errors than subjects from the Naïve Observer treatment, which is consistent with H2. However, the data also reveals that there is very little difference in the LE levels between the two treatments which is not consistent with H2. Taken together, the experimental results suggest that LBO can be, perhaps, marginally more effective when observing an optimal decision making opponent.

Table 5
Error Rate Comparison Between Optimal Observer and Naïve Observer

	Per Player Mean Error Rate	
Measure	Optimal Observer	Naïve Observer
E-3	35%**	49%
E-6	40%	41%
E-9	47%	42%
E-12	40%**	57%
E-15	32%	31%
E-18	19%	9%
Any Error	35%	38%

Notes: All error rate measures were tested using a 1-sided Mann-Whitney test
* significantly smaller at the 10% level ** significantly smaller at the 5% level

Table 6
LE Comparison Between Optimal Observer and Naïve Observer

	Per Player Mean LE	
Round	Optimal Observer	Naïve Observer
8	11.50	11.75
10	9.38	8.88
12	8.13	9.63
Aggregate LE	9.67	10.08

Notes: All LE measures were tested using a 1-sided Mann-Whitney test

3.3 Testing H3 and H4: LBD compared to LBO

To test H3 and H4, I compare the decision making between players and observers for both the *optimal opponent* and *naïve opponent*. Table 7 compares the aggregate proportions of MFGs per round, and the per subject mean total MFGs. For all rounds, I fail to reject the null hypothesis that the proportion of MFGs is equal for players and observers, for both the *optimal opponent* and *naïve opponent*, using a 2-sided Fisher's Exact test. Similarly, the mean proportion of total MFGs is not significantly different between players and observers, for either the *optimal opponent* or *naïve opponent*, using a 2-sided Mann-Whitney U-test.

TABLE 7
MFG Comparison between Players and Observers

	Aggregate Proportion of MFG			
	Optimal Opponent		*Naïve Opponent*	
Round	Player	Observer	Player	Observer
7	4/24	8/24	6/24	5/24
8	9/24	4/24	3/24	0/24
9	6/24	10/24	5/24	5/24
10	8/24	6/24	3/24	4/24
11	8/24	8/24	4/24	4/24
12	8/24	7/24	6/24	2/24
Total MFG / Player	1.79/6	1.79/6	1.13/6	.83/6

Notes: The proportions of MFGs were tested using 2-sided Fisher's Exact test.
The Total MFG/Player was tested using a 2-sided Mann-Whitney test.

Table 8 and Table 9 compares the mean error rates and mean LE, respectively. For both the *optimal opponent* and *naïve opponent,* and for all error rates considered (with the exception of E-9 for the *optimal opponent*), I fail to reject the null hypothesis that the mean error rates are equal between players and observers. Table 9 reveals a similar pattern in the data for LE. There are no significant differences in any LE measure between players and observers, for either the *optimal opponent* or *naïve opponent.*

Overall, the experimental data is generally consistent with H3 and H4. That is, players and observers appear to exhibit comparable decision making quality regardless of their opponent's decision making quality. Specifically, for both the *optimizing opponent* and *naïve opponent*, there are no significant differences in MFGs, error rates, or LE measures between the players and observers. The experimental results suggest that LBD and LBO are comparable in terms of their effectiveness at improving strategic decision making in G21; This seems to be true irrespective of the decision making quality exhibited by one's opponent.

Table 8
Error Rate Comparison Between Players and Observers

	Per Player Mean Error Rate			
	Optimal Opponent		*Naïve Opponent*	
Measure	Player	Observer	Player	Observer
E-3	26%	35%	47%	49%
E-6	41%	40%	43%	41%
E-9	28%**	47%	40%	42%
E-12	36%	40%	50%	57%
E-15	23%	32%	33%	31%
E-18	16%	19%	11%	9%
Aggregate Error Rate	29%	35%	36%	38%

Notes: All error rate measures were tested using a 2-sided Mann-Whitney test
** significantly different at the 5% level

Table 9
LE Comparison Between Players and Observers

	Per Player Mean LE			
	Optimal Opponent		*Naïve Opponent*	
Round	Player	Observer	Player	Observer
8	7.38	11.50	9.38	11.75
10	6.50	9.38	9.50	8.88
12	6.25	8.13	7.25	9.63
Aggregate Mean LE	6.71	9.67	8.71	10.08

Notes: All LE measures were tested using a 2-sided Mann-Whitney test

4 Conclusion

The motivation of this paper was twofold: The first was to experimentally test whether LBD and LBO are more effective in strategic settings when agents face an opponent who exhibits optimal decision making, relative to sub-optimal decision making. The second was to experimentally compare the effectiveness of LBD with LBO in a strategic setting for both an optimal and sub-optimal decision making opponent. To shed light on these research questions, I considered an experimental design that featured a stylized 2-player, dynamic game, and the implementation of optimal and naïve decision making computer opponents.

I find that subjects who initially played G21 against the *optimizing opponent* exhibited significantly better decision making than those subjects who initially played G21 against the *naïve opponent.* While, subjects who initially observed the *optimizing opponent* exhibited, at most, marginally better decision making to those subjects who initially observed the *naïve opponent.* Taken together, the experimental results provide evidence that LBD can be significantly more effective in strategic settings when subjects face an optimal decision making opponent relative to a naïve decision making opponent. Whereas, LBO may only be marginally more effective when observing an optimal decision making opponent relative to a naïve decision making opponent.

Additionally, I find that conditional on the decision making quality of the opponent, there is little difference in decision making quality exhibited by subjects who initially played and subjects who initially observed. That is, LBD and LBO appear to be comparably effective regardless of the decision making quality of the opponent. By comparison, Merlo and Schotter (2003) find experimental evidence that LBO is more effective than LBD. However, Merlo and Schotter compare LBD and LBO using a single-agent, single-decision maximization problem, while I compare LBD and LBO using a 2-player dynamic game. Therefore, results of this paper should not be viewed in opposition to those of Merlo and Schotter, but rather as complementary. Investigating which settings are more conducive to LBD and LBO, and the reasons why, remain open questions for future research.

Before I conclude, I speculate about possible mechanisms that could explain why LBD is more effective in strategic settings when agents play an optimal decision making opponent. One possibility is some type of reinforcement learning (Erev and Roth (1998), and Camerer and Ho (1999)). Subjects playing against the *optimizing opponent* who fail to play the dominant solution lose the game; hence, they negatively reinforce the non-dominant strategies and gradually learn to play a dominant strategy. Another possibility is that subjects who play against the *optimizing opponent*, in lieu of the *naïve opponent*, are simply "mimicking" the *optimizing opponent.*[19] However,

[19] Vega-Redondo (1997) and Schlag (1999) develop formal models of imitation learning. However, both models apply to simultaneous move games where the strategies are perfectly observable. G21

the data provides some evidence of gradual learning in G21, within a given game and across games, which is less consistent with simply mimicking the dominant solution. Identifying why LBD is more effective when the opponent makes better decisions could be an area for future research.

Although the strategic setting considered in this study was abstract and atypical of many strategic settings in the field, a few important insights can be gleaned. First, LBO can possibly be just as effective as LBD in complex strategic settings. This is particularly relevant in settings where LBD is costly, as it may behoove an agent to initially gather knowledge by observing in lieu of doing. The results from this paper suggest that apprenticing first can be just as effective as jumping right in, and potentially less costly. Second, LBD can be more effective in strategic settings when facing an opponent who makes good decisions. This suggests that agents who are motivated to become better decision makers as effectively as possible, through experience, should gather that experience by playing against good decision making opponents.

is a sequential game and, therefore, the complete strategy is not observed, only actions. Akin to DSB (2010), I refer to this type learning in G21 as "mimicking" and not imitation learning.

5 Appendix

5.1 Experimental Instruction for PLAYER

This is an experiment in strategic decision making. Please read these instructions carefully and pay attention to any additional instructions given by the experimenter. You have the potential to earn additional compensation and the amount depends on the decisions that you make in this experiment.

You have been assigned the PLAYER role. You will be playing 12 rounds of a simple two player game that will be described below. You will be playing against a programmed computer opponent for all 12 rounds. In addition, for the first 6 rounds, an observer will be standing behind you watching you play against the computer. You are not permitted to communicate with the observer in any capacity. After 6 rounds, the observer will be instructed to move to an empty computer carol and you will be instructed to play 6 more rounds.

After you finish playing all 12 rounds, please remain seated and an experimenter will come by and pay you your additional compensation. In addition to the $5 show-up fee, you will be compensated $1 for each round that you win. The observer will also be compensated $1 for each round that you win during the first 6 rounds. After you have been paid, please quietly leave the laboratory. If you have any questions at any time, raise your hand and an experimenter will be by to answer your question.

5.2 Experimental Instruction for OBSERVER

This is an experiment in strategic decision making. Please read these instructions carefully and pay attention to any additional instructions given by the experimenter. You have the potential to earn additional compensation and the amount depends on the decisions that you make in this experiment.

You have been assigned the OBSERVER role. You will first be watching 6 rounds, and then be playing 6 rounds of a simple two player game that will be described below. For the first 6 rounds, you will be standing behind a player watching him/her play against the programmed computer opponent. You are not permitted to communicate with the player in any capacity. After 6 rounds, you will be instructed to move to an empty computer carol where you will then play 6 rounds of the game against a programmed computer opponent.

After you finish playing the last 6 rounds, please remain seated and an experimenter will come by and pay you your additional compensation. In addition to the $5 show-up fee, you will be compensated $1 for each round that you win while you are playing, and $1 for each round that the player wins while you are observing. After you have been paid, please quietly leave the laboratory. If you have any questions at any time, raise your hand and an experimenter will be by to answer your question.

5.3 Sample Screen Shot of Instructions

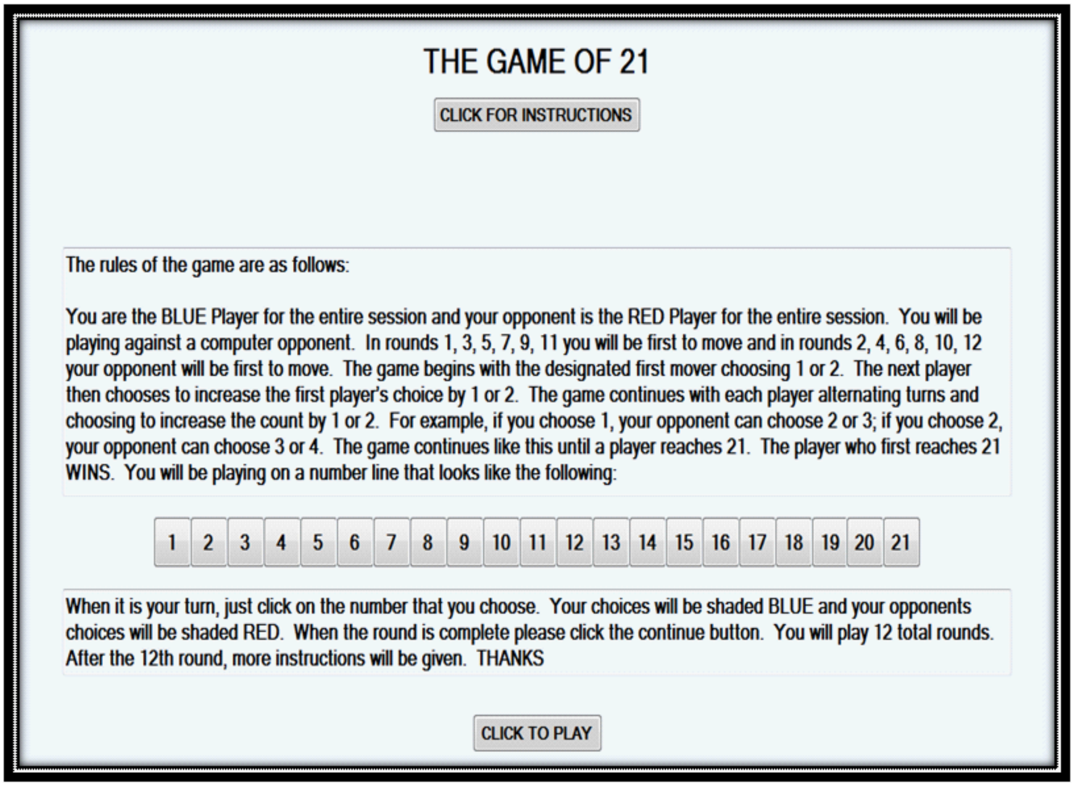

5.4 Sample Screen Shot of G21 Interface

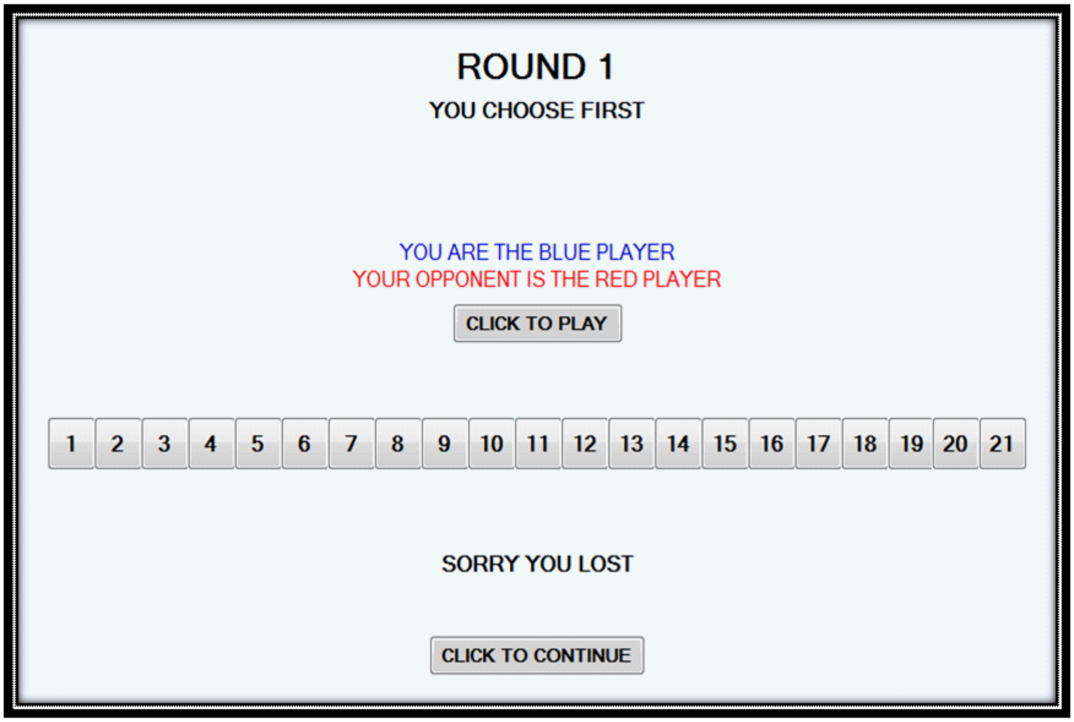

References

[1] Arrow, K., 1962. The Economic Implications of Learning by Doing. The Review of Economic Studies 29, 155-173.

[2] Bikhchandani, S., Hirshleifer, D., Welch, I., 1998. Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. The Journal of Economic Perspectives 12, 151-170.

[3] Binmore, K., Shaked, A., Sutton., 1985. Testing No cooperative Bargaining Theory: A Preliminary Study. American Economic Review 75, 1178-1180.

[4] Binmore, K., McCarthy, J., Ponti, G., Samuelson, L., Shaked, A., 2002. A Backward Induction Experiment. Journal of Economic Theory 104, 48-88.

[5] Bosch-Domenech, A., Montalvo, J., Nagel, R., Satorra, A., 2002. One, Two, (Three), Infinity, . . . : Newspaper and Lab Beauty Contest Experiments. American Economic Review 92, 1687-1701.

[6] Camerer, C., Ho, T-H., 1999. Experienced-Weighted Attraction Learning in Normal Form Games. Econometrica 67, 827-874.

[7] Carpenter, J., 2003. Bargaining Outcomes as a Result of Coordinated Expectations: An Experimental Study of Sequential Bargaining. Journal of Conflict Resolution 47, 119-141.

[8] Chou, E., McConnell, M., Nagel, R., Plott, C., 2009. The control of Game Form Recognition in Experiments: Understanding Dominant Strategy Failures in a Simple Two Person 'Guessing' Game. Experimental Economics 12, 159-179.

[9] Dixit, A. 2005. Restoring Fun to Game Theory. Journal of Economic Education 36, 205-219.

[10] Duffy, J., Nagel, R., 1997. On the Robustness of Behavior in Experimental 'Beauty Contest' Games. The Economic Journal 107, 1684-1700.

[11] Dufwenberg, M., Sundaram, R., Butler, D., 2010. Epiphany in the Game of 21. Journal of Economic Behavior and Organization.

[12] Dursch, P., Kolb, A., Oechssler, J., Schipper, B., 2010. Rage Against the Machines: How Subjects Learn to Play Against Computers. Economic Theory 43, 407-430.

[13] Erev, I., Roth, A., 1998. Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibrium. The American Economic Review 88, 848-881.

[14] Fey, M., McKelvey, R., Palfrey, T., 1996. An Experimental Study of Constant-Sum Centipede Games. International Journal of Game Theory 25, 269-289.

[15] Gneezy, U., Rustichini, A., Vostroknutov, A., 2010. Experience and Insight in the Race Game. Journal of Economic Behavior and Organization 75, 144-155.

[16] Grosskopf, B., Nagel, R., 2008. The Two Person Beauty Contest. Games and Economic Behavior 62, 93-99.

[17] Hall, K. 2009. A Pure Test of Backward Induction. Doctoral Dissertation: University of Tennessee.

[18] Harrison, G., McCabe, K., 1992. The Role of Experience for Testing Bargaining Theory in Experiments. In: Isaac, M. (Ed.). Research in Experimental Economics, Vol. 5. Greenwich: JAI Press, 137-169.

[19] Harrison, G., McCabe, K., 1996. Expectations and Fairness in a Simple Bargaining Experiment. International Journal of Game Theory 25, 303-327.

[20] Hastie, R., 1986. Review essay: experimental evidence on group accuracy. In: Owen, G., Grofman, B. (Eds.). Information Pooling and Group Decision Making. Westport, CT: JAI Press, 129-57.

[21] Hill, G., 1982. Group versus individual performance: Are N + 1 heads better than one. Psychological Bulletin 91, 517-539.

[22] Ho, T-H., Camerer, C., Weigelt, K., 1998. Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty-Contests'. American Economic Review 88, 947-969.

[23] Jovanovic, B., Nyarko, Y., 1995. The Transfer of Human Capital. Journal of Economic Dynamics and Control 19, 1033–1064.

[24] John, E.R., Chesler, P., Bartlett, F., Victor, I., 1969. Observational learning in cats. Science 166, 901–903.

[25] Johnson, E., Camerer, C., Sen, S., Rymon, T., 2002. Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining. Journal of Economic Theory 104, 16-47.

[26] Laughlin, P., 1980. Social combination processes of cooperative problem-solving groups on verbal intellective tasks. In: Fishbein, M. (Ed.). Progress in Social Psychology. Hillsdale, NJ: Erlbaum, 127-155.

[27] Laughlin, P., Adamopoulos, J., 1980. Social combination processes and individual learning for six-person cooperative groups on an intellective task. Journal of Personality and Social Psychology 38, 941-947.

[28] Laughlin, P., Bonner, B., Miner, A., 2002. Groups perform better than the best individuals on letters-to-numbers problems. Organizational Behavior and Human Decision Processes 88, 605-620.

[29] Laughlin, P., Ellis, A., 1986. Demonstrability and social combination processes on mathematical intellective tasks. Journal of Experimental Social Psychology 22, 177-189.

[30] Levine, J., Moreland, R., 1998. Small groups. In: Gilbert D., Fiske, S., Lindzey, G. (Eds.). The Handbook of Social Psychology 4th Edition, Vol. 2. New York: McGraw-Hill, 415-69.

[31] Levitt, S., List, J., 2007. What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World. Journal of Economic Perspectives 21, 153-174.

[32] Levitt, S., List, J., Sadoff, S., Checkmate: Exploring Backward Induction Among Chess Players. Forthcoming in American Economic Review.

[33] McKelvey, R., Palfrey, T., 1992. An Experimental Study of the Centipede Game. Econometrica 60, 803-836.

[34] Merlo, A., Schotter, A., 2003. Learning by not Doing: An Experimental Investigation of Observational Learning. Games and Economic Behavior 42, 116-136.

[35] Nadler, J., Thompson, L., Van Boven, L., 2003. Learning Negotiation Skills: Four Models of Knowledge Creation and Transfer. Management Science 49, 529-540.

[36] Nagel, R., 1995. Unraveling in Guessing Games: An Experimental Study. American Economic Review 85, 1313-1326.

[37] Ochs, J., & Roth, A., 1998. An Experimental Study of Sequential Bargaining. American Economic Review 79, 355-384.

[38] Palacios-Huerta, I., & Volij, O., 2009. Field Centipedes. The American Economic Review 99, 1619-1635.

[39] Rapoport, A., Stein, W., Parco, J., Nicholas, T., 2003. Equilibrium Play and Adaptive Learning in a Three-Person Centipede Game. Games and Economic Behavior 43, 239-265.

[40] Rosenthal, R., 1981. Games of Perfect Information, Predatory Pricing, and the Chain-Store Paradox. Journal of Economic Theory 25, 92-100.

[41] Schlag, K., 1999. Which One Should I imitate? Journal of Math and Economics 31, 493-522.

[42] Shachat, J., Swarthout, T., 2004. Do we detect and exploit mixed strategy play by opponents. Mathematical Methods of Operations Research 59, 359-373.

[43] Siegel, S., & Castellan, Jr., N.J., 1988. Nonparametric Statistics for the Behavioral Sciences 2nd Edition. New York: McGraw-Hill Book Company.

[44] Thompson, P., 2010. Learning by Doing. In: Hall, B., Rosenberg, N. (Eds.). Handbook of the Economics of Innovation, Vol. 1. Amsterdam: Elsevier/North-Holland, 429-476.

[45] Tomasello, M., Davis-Dasilva, M., Camak, L., & Bard, K. (1987). “Observational Learning of Tool-Use by Young Chimpanzees,” Human Evolution 2, 175-183.

[46] Terkel, J., 1996. Cultural transmission of feeding behavior in black rats (rattus rattus). In: Heynes, C., Galef Jr., B. (Eds.). Social Learning in Animals: The Roots of Culture. San Diego: Academic Press, 17-45.

[47] Vega-Redondo, F., 1997. The Evolution of Walrasian Behavior. Econometrica 65, 375-384.