G Model
JEBO-2931; No. of Pages 11

ARTICLE IN PRESS

Journal of Economic Behavior & Organization xxx (2012) xxx–xxx

Contents lists available at SciVerse ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Learning to make better strategic decisions☆

Eric Cardella[a,b]

[a] Department of Economics, University of Arizona, McClelland Hall 401, PO Box 210108, Tucson, AZ 85721-0108, United States
[b] Department of Economics, Rochester Institute of Technology, 92 Lomb Memorial Drive, Rochester, NY 14623-5604, United States

ARTICLE INFO

Article history:
Received 7 May 2011
Received in revised form 15 March 2012
Accepted 18 April 2012
Available online xxx

JEL classification:
C73
C91
D83

Keywords:
Learning
Learning-by-doing
Learning-by-observing
Strategic game
Dominant strategy

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 2-player strategic game. Specifically, does LBD become more effective when agents face an opponent who exhibits optimal decision making? Similarly, does LBO become more effective when agents observe an opponent who exhibits 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. Whereas, LBO is, at most, marginally more effective when observing an optimal decision making opponent. The results also suggest that LBD is at least as effective as LBO at improving decision making in the 2-player game considered.

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 strategic decisions. Agents who lack high levels of strategic sophistication and/or deep reasoning ability are likely to initially 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 motivation of this paper is to investigate learning in strategic settings and provide insights regarding how agents can possibly become better strategic decision makers.

In relation to single agent decision tasks, one learning mechanism that can facilitate improved decision making is learning-by-doing (LBD). By repeatedly *doing* a decision task, an agent can acquire knowledge and skills that can subsequently lead to better decision making. 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). I refer the reader to Thompson (2010) for a comprehensive review of the extensive literature on LBD, including theoretical applications and empirical investigations supporting LBD. An alternative, yet related, learning mechanism that can facilitate improved decision making in single agent decision tasks is learning-by-observing (LBO). By repeatedly *observing* the decision making

☆ 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, Jörg Oechssler, John Wooders, an anonymous referee, and conference participants at the University of Arizona for additional helpful comments. I am grateful to the Economic Science Laboratory at the University of Arizona for providing financial support.

E-mail address: ecardell@email.arizona.edu

0167-2681/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jebo.2012.04.011

of another, an agent can acquire knowledge and skills 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]

In relation to strategic games, I contend that LBD corresponds to the acquisition of knowledge and skill by repeatedly playing the game. I, henceforth, refer to this analog of LBD in a game as *strategic* LBD. Similarly, LBO in a strategic game corresponds to the acquisition of knowledge and skill by repeatedly observing another agent play the game. I, henceforth, refer to this analog of LBO in a game as *strategic* LBO. Games, unlike single agent decision tasks, involve decision making of multiple agents. Therefore, it is possible that the effectiveness of strategic LBD and strategic LBO will be influenced by the decision making of the other agents in the game. The first motivation of this study is to experimentally investigate how strategic LBD and strategic LBO are influenced by the decision making quality of an agent's opponent. Specifically, I investigate whether strategic LBD becomes more effective when an agent repeatedly plays against an opponent who makes optimal decisions, compared to sub-optimal decisions. Similarly, I investigate whether strategic LBO becomes more effective when an agent repeatedly observes an agent who plays an opponent who makes optimal decisions, compared to sub-optimal decisions.

To shed light on these questions, I propose a stylized experimental design, described in detail in the following section, that uses a 2-player, sequential-move game which features a dominant strategy. The dominant strategy of the chosen game serves as an identifiable and measurable proxy for optimal strategic 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] In particular, 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. In the experiment, some subjects exclusively play the game, while other subjects initially observe a subject playing the game and then play the game themselves. This variation in whether subjects initially play the game or observe play of the game, in combination with the variation of the decision making quality of the computer opponent, allows me to identify how strategic LBD and strategic LBO are impacted by the decision making quality of one's opponent.

I find that subjects who initially played against the *optimizing opponent* make better decisions than subjects who initially played against the *naïve opponent*. However, I find that subjects who initially observed another subject playing the *optimizing opponent* make only marginally better decisions than those subjects who initially observed another subject playing the *naïve opponent*. These results suggest that strategic LBD can be more effective when playing against an opponent who makes an optimal decisions, while strategic LBO may be, at most, marginally more effective when observing an opponent who makes an optimal decisions.

As a second motivation of this study, I experimentally compare the effectiveness of strategic LBD and strategic LBO. That is, I compare the decision making quality of subjects who initially play the game to subjects who initially observe another subject play the game, for both the *optimizing opponent* and *naïve opponent*. This is similar in spirt to Merlo and Schotter (2003) who experimentally compare LBD and LBO in a single-agent profit maximization problem.[3] In their setting, Merlo and Schotter find that subjects who initially observe learn better than subjects who initially do. In the strategic game that I consider, I find very little difference between the decision making quality of the subjects who initially observe and the subjects who initially play. The results suggest that strategic LBD and strategic LBO appear to be comparably effective mechanisms for making better decisions. Because this study compares LBD and LBO in a strategic setting, the results should be viewed as complementary to those of Merlo and Schotter.

It is important to note that in a game with multiple decision makers, agents who play the game are going to also observe the decisions made by the other agents playing the game, i.e., observation of one's opponent is naturally embedded into playing a multi-decision maker game. In this regard, strategic LBD includes the effect of simultaneously observing the decision making of one's opponent while playing. Therefore, when I investigate strategic LBD, I am actually investigating the compound effect of playing the game *and* observing the opponent.[4] Because subjects are playing a game, it impossible to isolate the effect of playing the opponent, from observing the opponent. At the same time, because subjects who play necessarily observe the opponent, I contend that this compound effect is the appropriate and meaningful effect when investigating strategic LBD. However, to minimize the saliency of observing the opponent, with respect to strategic LBD, I consider two auxiliary treatments as part of the experimental design where subjects play an asymmetric version of G21. That is, a version of G21 where the subject and the computer opponent have different optimal strategies. In this asymmetric version of G21, I find

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

[2] The use of pre-programmed computer opponents is certainly not novel to this study. For example, 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. Dürsch 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 always make the same pre-specified decision.

[4] I thank an anonymous referee for calling to my attention this important point.

some evidence, although not as conclusive, that subjects who initially play against the *optimizing opponent* make marginally better decisions than subjects who initially play against the *naïve opponent*.

Ideally, one might want to gain insights regarding how agents learn to make better strategic decisions by experimenting in the field. However, identifying and measuring strategic decision making quality in the field can be difficult. Optimal decisions often depend on the agent's preferences and/or beliefs, both of which are difficult to measure. Without this information, 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 the field. By considering an experimental game in the lab that features a dominant strategy, I am able to identify and measure optimal strategic decision making. Additionally, the lab allows for the possibility to systematically control the decision quality of an agent's opponent by using pre-programmed computer opponents. Therefore, a stylized lab experiment enables me to gain insights regarding the impact of the opponents decision making quality on strategic LBD and strategic LBO, which would otherwise be difficult by experimentation or observation in the field. 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 chosen experimental game features a dominant strategy, which makes it atypical of many strategic settings in the field. Games that feature a dominant strategy are similar to intellective task, in the sense that there exists a demonstrably correct solution (Laughlin, 1980; Laughlin and Adamopoulos, 1980; Laughlin and Ellis, 1986). This similarity is relevant in light of several studies that have shown differences in learning between intellective tasks and judgement tasks (Hill, 1982; Hastie, 1986; Levine and Moreland, 1998; Laughlin et al., 2002 for a thorough review). I acknowledge that using a game featuring a dominant strategy may limit the generalizability (Levitt and List, 2007) of the insights gleaned from the experimental results regarding learning in a general strategic setting. Nevertheless, even at its most limited scope, this study can help better our understanding of LBD and LBO in strategic settings that feature dominant strategies. Such insights can be valuable when applied to games in the lab that feature dominant strategies, e.g., 2-player guessing games (Grosskopf and Nagel, 2008; Chou et al., 2009), or games in the field that feature dominant strategies, e.g., second-price auctions.[5]

The paper proceeds by formally describing the experimental design and developing the research hypotheses in Section 2. The results are presented and discussed in Section 3, and Section 4 concludes.

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.[6]

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 a multiple of three, when it is available, is dominant for the second mover.[7] 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 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 features a binary outcome of "win" or "lose". This eliminates possible

[5] 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).

[6] 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. (2011) 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.

[7] 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.

Please cite this article in press as: Cardella, E., Learning to make better strategic decisions. J. Econ. Behav. Organ. (2012), http://dx.doi.org/10.1016/j.jebo.2012.04.011

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.[8]

In order to investigate strategic LBD and strategic LBO, it is crucial that the chosen game be complex enough to allow for learning. Hence, it must be the case that most subjects playing G21 for the first time fail to play the dominant solution, but exhibit better decision making over time. This is confirmed by DSB (2010) who find that roughly 85 percent of subjects playing G21 for the first time initially fail to play the dominant solution. Similarly, Gneezy et al. (2010) 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. These 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, which enable me to better quantify deviations from optimal decision making.

G21 is a stylized strategic game and this 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 strategic LBD and strategic LBO.[9]

2.2. Experimental treatments and research hypotheses

There were two possible types of computer opponents in the experiment, 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.[10]

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

Optimal Player Treatment: The subject played six rounds of G21 against the *optimizing opponent* followed by six rounds of G21 against the *naïve opponent*.
Naïve Player Treatment: The subject played all twelve rounds of G21 against the *naïve opponent*.
Optimal Observer Treatment: 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 Treatment: 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*.

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 and, 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, which isolates the treatment effect.[11] Second, playing the final six rounds

[8] 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.

[9] The use of 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.

[10] Hall (2009) implements a similar design where subjects play variations of the race game against a pre-programmed computer opponent. 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 *optimizing opponent*. Although the design in this paper features some similar aspects to that of Hall (the use of the race game and the use of an optimizing pre-programmed computer opponent), the research questions differ substantially. Namely, Hall experimentally investigates how mid-game *teaser* payments, both on and off the Backward Induction path, affect a subject's ability to learn the Backward Induction solution, while this paper investigates how the decision making quality of one's opponent affects strategic LBD and strategic LBO.

[11] Technically, there is an additional implicit difference between the treatments besides the main treatment effect. In particular, in both the Optimal Player Treatment and Optimal Observer Treatment, the opponent switched from the naïve opponent to the optimizing opponent after the first six rounds. It should be noted that this change in the opponent quality could have, in theory, had an effect on the behavior of subjects in these treatments that was distinct from the subject's main treatment effect in the first six rounds. However, there are several reasons why this possible secondary affect is likely to have had a negligible impact on subjects' decision making. First, subjects in these two treatments were unaware that the opponent was even changing after six rounds and it seems unlikely that they anticipated or recognized such a change. Second, the decision making of the opponent has no impact on the optimal strategy of choosing every available multiple of three in the game. Hence, under the reasonable assumption that subjects are trying to win every game, changing the decision making quality of the opponent should not have directly influence the strategy of subjects. I thank the referee for pointing out this additional difference and providing value insights about why it is likely not an issue.

Please cite this article in press as: Cardella, E., Learning to make better strategic decisions. J. Econ. Behav. Organ. (2012), http://dx.doi.org/10.1016/j.jebo.2012.04.011

against the *naïve opponent* allows me to construct of a measure of decision making quality (defined and described in the next section) that captures and quantifies the *extent* to which subjects are failing to make optimal decisions in G21, which will provide a more robust conclusion about how the decision making quality of their opponent impacts strategic LBD and strategic LBO.[12]

Recall, that the first motivation of this study is to investigate whether (i) strategic LBD and (ii) strategic LBO are more effective when facing an optimal decision making opponent, compared to a sub-optimal decision making opponent. To answer these questions, I use 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, and (H2) subjects from the Optimal Observer Treatment exhibit better decision making than subjects from the Naïve Observer Treatment.

The second motivation of this study is to compare the effectiveness of strategic LBD and strategic LBO, for both the (iii) *optimizing opponent* and the (iv)*naïve opponent*. I use the experimental data to compare of strategic LBD and strategic 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, and (H4) subjects from the Naïve Player Treatment exhibit equivalent decision making quality as subjects from the Naïve Observer Treatment.

2.3. Experimental procedure

The subject pool for this experiment consisted of undergraduates from the University of Arizona, and all of the sessions were conducted in the Economic Science Laboratory (ESL). 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. All subjects were informed that they were playing against a computer opponent, but received no information about the strategy implemented by the computer.[13]

Twenty-four subjects participated in each of the four treatments. In each experimental session, 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 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 designated computer opponent, while the observers merely observed the play.[14] After observing the first 6 rounds of the game, 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. Each subject began as the first mover and alternated between first mover and second mover in all subsequent rounds.

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

3. Results

Testing H1–H4 requires measuring decision making quality in G21. To help minimize measurement error and provide a more well-rounded and robust test of these research hypotheses, I consider two complementary measures that characterize decision making quality in G21. I compare the treatments using the last six rounds of the game when all subjects played G21 against the *naïve opponent*. Further, all hypotheses are tested using subject level aggregate data from the three rounds when subjects acted as the second mover and, thus, had an opportunity to play the dominant solution. This is intended to

[12] This is because in rounds where subjects play the *optimizing opponent* very little insight can be gained regarding a subject's decision making quality beyond whether they played the optimal strategy of choosing every multiple of three. However, when subjects play against the *naïve opponent*, opportunities will likely arise at various subgames, to begin playing a dominant strategy. Because the *naïve opponent* randomizes equally at every decision node, if a subject fails to choose an available multiple of three, then there is a 50 percent chance that a multiple of three will be available on the subjects next turn.

[13] In particular, subjects were not informed about whether they were facing the *optimizing opponent* or the *naïve opponent*. This was intended to help isolate the effect of the opponent's decision making quality on strategic LBD and strategic LBO by eliminating any possible informational effects that could arise by informing subjects about the decision quality of the opponent. It is certainly possible that providing information about the decision making quality of the opponent could affect the results, as noted by an anonymous referee. However, absence of such information does not undermine or nullify the research hypotheses regarding how strategic LBD and strategic LBO are impacted by the actual decision quality of the opponent. Investigating how information regarding the opponent' decision quality affects strategic LBD and strategic LBO could be an interesting area of future research, although beyond the scope of the current study.

[14] 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.

Table 1
Comparison of Naïve Player data for first half and second half – G21.

Measure	Naïve Player Treatment		(*p*-Value)
	First half (rounds 2, 4, 6)	Second half (rounds 8, 10, 12)	
Perfect Games	.25/3 (8%)	.50/3 (17%)	(0.070)
Total Error Rate	45%	38%	(0.042)

Notes: Both measures are reported as subject level averages, and were tested using a 1-sided Wilcoxon matched-pairs signed-ranks test. Significance levels are robust to a 1-sided matched pairs sign-test.

help reduce bias in measurement of decision making quality, smooth noise in individual rounds of the game, and ensure observational independence across treatments.

The first measure of decision making quality I consider is a *Perfect Game*. I define a Perfect Game as a game of G21 where the subject played the dominant solution, i.e. chose 3, 6, 9, 12, 15, 18, and 21.[15] Playing the dominant solution of choosing all multiples of three in G21 guarantees victory. Hence, I assume that playing a Perfect Game serves as a proxy measure for optimal decision making in G21, and playing more Perfect Games can be viewed as evidence consistent with better decision making in G21.[16] I point out that the Perfect Game measure is only a proxy for optimal decision making. The reason being is that when a subject plays against the *naïve opponent*, the subject can fail to play a Perfect Game and still win. However, if a subject plays a Perfect Game, then the win probability is equal to one. Whereas, if a subject does not play a Perfect Game, then there is a chance that a subject could lose the game and the win probability is strictly less than one. Therefore, in expectation, playing a Perfect Game yields a higher payoff, which is why I consider the Perfect Game measure a suitable proxy for optimal decision making in G21.

However, a subject may fail to play a Perfect Game, but nevertheless, exhibit comparably better decision making in G21. For example, it seems reasonable to assume that a subject who failed to choose 3, but proceeded to choose every subsequent available multiple of three in the game, exhibited better decision making than a subject who failed to choose every available multiples of three in a game. In order to compare decision making quality in non Perfect Games, I consider an additional measure of decision making quality in G21. The second measure is *Total Error Rate*, which I define as the total proportion of all available multiples of three that a subject failed to choose. Because a failure to choose an available multiple of three is sub-optimal, a lower Total Error Rate can viewed as evidence of better decision making in G21. The Total Error Rate measure is intended to serve as an overall "birds eye view" measure of decision making quality in G21.

3.1. Learning in G21

Testing H1–H4 are only possible conditional on the maintained assumptions that G21 is sufficiently complex that: (A1) most subjects initially fail to make optimal decision, and (A2) sufficiently straightforward that learning is possible. Hence, the ability to test how strategic LBD and strategic LBO are affected by the decision making quality of a subject's opponent. To verify A1, I look at the data from round 2, which is the first round when subjects acted as the second mover and, thus, had an opportunity to play a Perfect Game. In round 2, only 7/48 (15 percent) subjects, aggregated over the Optimal Player and Naïve Player Treatments, played a Perfect Game of G21. I contend that the 85 percent failure rate, which is similar to that found in DSB (2010), is sufficiently high to support the claim that most subjects initially fail to play G21 optimally and, thus, A1 is verified.

To verify A2, I compare the decision making of subjects in the Naïve Player Treatment from the first half (rounds 2, 4, 6) with the second half (rounds 8, 10, 12). Table 1 reveals that subjects in the Naïve Player Treatment played significantly more Perfect Games and had significantly lower Total Error Rates in the second half compared to the first half. Hence, subjects exhibited better decision making in the second half compared to the first half, which provides evidence that learning is possible in G21 and verifies A2. Given A1 and A2 are verified, it is possible to proceed with testing H1–H4.

3.2. Hypotheses testing

I first test H1, namely, subjects from the Optimal Player Treatment exhibit better decision making than subjects from the Naïve Player Treatment. Table 2 reports the aggregate comparison of Perfect Games and Total Error Rates. From Table 2 we see that the average proportion of Perfect Games is 1.04/3 (35 percent) in the Optimal Player Treatment compared to.50/3 (17 percent) in the Naïve Player Treatment, which is narrowly insignificant ($p = 0.106$). The average Total Error Rate is 25 percent for subjects in the Optimal Player Treatment compared to 38 percent for the Naïve Player Treatment, which is significant ($p = 0.007$). Hence, both measures yield differences that are in the decision of better decision making for subjects

[15] A Perfect Game is analogous to the measure of "perfect play" considered by DSB (2010).

[16] Because playing an Perfect Game requires choosing seven sequential multiples of three, it is unlikely that a subject would just randomly play a Perfect Game. The probability of playing a Perfect Game 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 Perfect Game is different across treatments.

Please cite this article in press as: Cardella, E., Learning to make better strategic decisions. J. Econ. Behav. Organ. (2012), http://dx.doi.org/10.1016/j.jebo.2012.04.011

Table 2
Effectiveness of strategic LBD and strategic LBO – G21.

Measure	Player subjects			Observer subjects		
	Optimal Player Treatment	Naïve Player Treatment	(*p*-Value)	Optimal Observer Treatment	Naïve Observer Treatment	(*p*-Value)
Perfect Games	1.04/3 (35%)	.50/3 (17%)	(0.106)	.71/3 (24%)	.25/3 (8%)	(0.046)
Total Error Rate	25%	38%	(0.007)	36%	41%	(0.252)

Notes: Both measures are reported as subject level averages, and were tested using a 1-sided Mann–Whitney *U*-test.

Table 3
Comparison of strategic LBD with strategic LBO – G21.

Measure	Optimizing opponent			Naïve opponent		
	Optimal Player Treatment	Optimal Observer Treatment	(*p*-Value)	Naïve Player Treatment	Naïve Observer Treatment	(*p*-Value)
Perfect Games	1.04/3 (35%)	.71/3 (24%)	(0.601)	.50/3 (17%)	.25/3 (8%)	(0.430)
Total Error Rate	25%	36%	(0.108)	38%	41%	(0.934)

Notes: Both measures are reported as subject level averages, and were tested using a 2-sided Mann–Whitney *U*-test.

in the Optimal Player Treatment, one of which is significant and the other only very narrowly misses being significant. Taken together, this provides evidence that is largely consistent with H1, which suggests that strategic LBD can be more effective when agents play an optimal decision making opponent, relative to a sub-optimal decision making opponent.

I proceed by testing H2, namely, subjects from the Optimal Observer Treatment exhibit better decision making than subjects from the Naïve Observer Treatment. Table 2 reports the aggregate comparison of Perfect Games and Total Error Rates. The average proportion of Perfect Games is.71/3 (24 percent) in the Optimal Observer Treatment compared to .25/3 (8 percent) in the Naïve Observer Treatment, which is significant ($p = 0.046$). Whereas, the average Total Error Rate is 36 percent for subjects in the Optimal Observer Treatment and 41 percent for the Naïve Observer Treatment, which is not significant ($p = 0.252$). Although both measures yield differences that are in the direction of better decision making for the Optimal Observer Treatment, only one of the measures is significant. Taken together, the data provides some evidence, although not conclusive, that is in the direction consistent with H2, which suggests that strategic LBO is, at most, marginally more effective when observing an optimal decision making opponent.[17]

To test H3 and H4, I compare the decision making between players and observers for both the *optimizing opponent* and *naïve opponent*. Table 3 reports the relevant comparisons of the Perfect Games and Total Error Rates. Looking first at when subjects faced the *optimizing opponent*, Table 3 reveals that the average proportion of Perfect Games is higher in the Optimal Player Treatment compared to the Optimal Observer Treatment, although not statistically significant ($p = 0.601$). Similarly, the average Total Error Rate is lower in the Optimal Player Treatment compared to the Optimal Observer Treatment, although narrowly insignificant ($p = 0.108$). Looking next at when subjects faced the *naïve opponent*, Table 4 reveals a similar pattern in Perfect Games and Total Error Rates between subjects in the Naïve Player and Naïve Observer Treatments, with neither of the two measures being statistically significantly different.

The data reveals no significant difference in decision making between player subjects and observer subjects when facing either the *optimizing opponent* or the *naïve opponent*, which is consistent with H3 and H4. Additionally, all of these insignificant differences are in the direction of better decision making for the player subjects compared to the observer subjects. Hence, the results suggest that strategic LBD appears to be no less effective than strategic LBO at improving decision making in G21, regardless of the decision making quality of the opponent.

3.3. *Results from an asymmetric version of G21*

I have found evidence that subjects from the Optimal Player Treatment exhibit better decision making in G21 than subjects from the Naïve Player Treatment. However, as I have previously acknowledged in the introduction, one possible issue is that player subjects also get to observe the opponent. That is, there is an observational component to strategic LBD in the form of observing the computer opponent. This observational component is particularly salient in G21 due to the symmetric structure of the game. The ability to observe the computer opponent, in combination with the symmetry of the

[17] In addition to the decision making quality of the computer opponent, it is possible that the effectiveness of strategic LBO might have also been affected by the decision making quality of the corresponding human player, as aptly noted by an anonymous referee. Unfortunately, the data does not allow me to uniquely match the observer subject with their corresponding player subject. Thus, investigating how the decision making quality of the player subject affects strategic LBO is not possible given the current data, although it certainly could provide interesting insights. However, because the decision making quality of the computer opponent had only a marginal effect of the observers, it is likely that the decision making quality of the human subject had, at most, a similarly marginally significant effect on decision making quality of the observer.

dominant solution between the player subject and computer opponent in G21, could be the primary reason for the enhanced decision making exhibited by subjects in the Optimal Player Treatment.[18]

To investigate this possibility further, I consider two auxiliary treatments where subjects repeatedly play an asymmetric version of G21. In this asymmetric version of G21, which I henceforth refer to as G21(A), the player subject must increment the count by 1 or 2, while the computer opponent must increment the count by 2 or 3. G21(A) features a first-mover advantage for both the player subject and the computer opponent. However, the asymmetry in increment length induces an asymmetric dominant solution. In particular, the dominant solution for the first-moving player subject is to choose 1, 5, 9, 13, 17, and 21, while the dominant solution for first-moving computer opponent is to choose 2, 6, 10, 14, 18, and 21. Again, I consider the same two types of pre-programmed computer opponents for G21(A); the *naïve opponent* that randomly increments the count by 2 or 3, and the *optimizing opponent* that plays the dominant solution.[19] Each of the two auxiliary treatments are described as follows:

Optimal Player(A) Treatment: The subject played ten rounds of G21(A) against the *optimizing opponent* followed by six rounds of G21(A) against the*naïve opponent*.[20]
Naïve Player(A) Treatment: The subject played all 16 rounds of G21(A) against the *naïve opponent.*

A total of 31 subjects participated in the Optimal Player(A) Treatment and 36 in the Naïve Player(A) Treatment. As with the previous treatments, subjects alternated as the first and second mover beginning as the first mover, and were paid using a similar scheme. These two additional treatments based off G21(A) allow me to re-test how the decision making quality of one's opponent impacts strategic LBD, while minimizing the saliency of the opponent observation component. That is, I re-test H1 by comparing the decision making quality from the Optimal Player(A) and Naïve Player(A) Treatments. I again use the Perfect Game and Total Error Rate measures, after making the appropriate adaptations to suit G21(A), to evaluate decision making quality.

I again need to verify (A1) that G21(A) is sufficiently complex that most subjects initially fail to make optimal decision, and (A2) sufficiently straightforward that learning is possible. The data reveals that only 3/67 (4 percent) of subjects play a Perfect Game in round 1, which verifies A1. In regards to A2, Table 4 compares the aggregate decision making data from the Naïve Player(A) Treatment for the first half (rounds 1, 3, 5, 7) and the second half (rounds 9, 11, 13, 15). From Table 2, we can see that, on average, subjects played significantly more Perfect Games and had a lower Total Error Rate in the second half rounds, which verifies A2. As a result of verifying A1 and A2, it is possible to proceed in re-testing H1.

Table 5 compares the decision making data from the Optimal Player(A) and Naïve Player(A) Treatments. From Table 5, we can see that the aggregate proportion of Perfect Games is.98/3 (33 percent) in the Optimal Player(A) Treatment compared to.69/3 (23 percent) in the Naïve Player(A) Treatment, which is narrowly insignificant ($p = 0.134$). The Total Error Rate is 25 percent in the Optimal Player(A) Treatment compared to 33 percent in the Naïve Player(A) Treatment, which is marginally significant ($p = 0.098$). Given these significance levels of ($p = 0.134$) and ($p = 0.098$) when comparing across treatments, the data does not provide conclusive evidence that subjects from the Optimal Player(A) Treatment make significantly better decisions that subjects from the Naïve Player(A) Treatment. However, both measures yield differences that are in the direction of better decision making for subjects in the Optimal Player(A) Treatment, one of which is significant and the other borders on being significant. Taken together, this provides some evidence that is consistent with H1, which suggests that, even in an asymmetric strategic setting where the opponent observation component is rendered less salient, strategic LBD can be marginally more effective when agents play an optimal decision making opponent.

3.4. Discussion

In this section, I provide some discussion about two plausible learning mechanisms that subjects could be implementing, in relation to the results from each of the treatments in G21 and G21(A). The first is subjects are implementing some kind of "imitation" or "mimicking" of the choices made by the opponent, if the opponent wins.[21,22]. Recall that the *optimizing opponent* would play the dominant solution of choosing every available multiple of three and, therefore, wins every time it

[18] I thank an anonymous referee for keenly calling attention to this point.

[19] Technically, the dominant solution for the first-moving computer opponent is to choose from the following sets: {2}, {5,6}, {9,10}, {13,14}, {17,18}, {21}. However, to induce as much possible asymmetry, the optimal computer opponent was programmed to choose 2, 6, 10, 14, 18, and 21 whenever possible. Hence, in games when the computer acted as the first mover, the optimizing opponent did play the 2, 6, 10, 14, 18, and 21 solution.

[20] Because G21(A) is more complex than G21, in the treatments where subjects played a total of 16 rounds instead of 12 rounds to ensure enough time for ample learning. However, the comparison of decision making across treatments will still be based off the last 6 rounds when all subjects played against the *naïve opponent*.

[21] Vega-Redondo (1997) and Schlag (1999) develop formal models of imitation learning where agents imitate the strategies of others who obtain higher payoffs. Both of these models are developed in a static framework. However, G21 is a sequential game and the complete strategy of the opponent is not observed upon completion of the round. Therefore, agents playing G21 would be unable to imitate the strategy of the opponent in G21 as specifically modelled by Vega-Redondo (1997) and Schlag (1999). In that regard, it may be more pedagogical to refer to this type of mechanism in G21 as mimicking, akin to the terminology used in DSB (2010).

[22] Experimental evidence consistent with imitation has been subsequently documented, e.g., Huck et al. (1999, 2000), Offerman et al. (2002), and Apesteguia et al. (2007).

Table 4
Comparison of Naïve Player data for first half and second half – G21A.

Measure	Naïve Player(A) Treatment		(*p*-Value)
	First half (rounds 1, 3, 5, 7)	Second half (rounds 9, 11, 13, 15)	
Perfect Games	.53/4 (13%)	.92/4 (23%)	(0.009)
Total Error Rate	42%	32%	(0.003)

Notes: Both measures are reported as subject level averages, and were tested using a 1-sided Wilcoxon matched-pairs signed-ranks test. Significance levels are robust to a 1-sided matched pairs sign-test.

Table 5
Comparison of Optimal Player(A) and Naïve Player(A) data – G21A.

Measure	Optimal Player(A) Treatment	Naïve Player(A) Treatment	(*p*-Value)
Perfect Games	.98/3 (33%)	.69/3 (23%)	(0.134)
Total Error Rate	25%	33%	(0.098)

Notes: Both measures are reported as subject level averages, and were tested using a 1-sided Mann–Whitney *U*-test.

acted as the second-mover or the subject failed to play the dominant solution. Hence, mimicking the *optimizing opponent* in G21 would result in choices more in line with the dominant solution of choosing every available multiple of three. As a result, if subjects were mimicking, then we would expect the decisions of subjects in the Optimal Player Treatment to be more in line with the dominant solution than in the, *Naïve Player Treatment* i.e., a higher proportion of Perfect Games and a lower Total Error Rate. This is largely consistent with what is observed in the data.

However, if mimicking was the *sole* mechanism by which subjects learned in G21, then we would expect very little difference in the propensity to exhibit better decision making between subjects who played against the *optimizing opponent* and the *naïve opponent* in G21(A), where mimicking *does not* result in choices that are more in line with the dominant solution. However, I find some evidence in the direction that suggests the decisions from subjects in the Optimal Player(A) Treatment were marginally more in line with the optimal solution than in the Naïve Player(A) Treatment, in the form of a lower Total Error Rate. This suggests that mimicking is likely not the sole learning mechanism being implemented.

The second plausible learning mechanism is some kind of reinforcement learning.[23] When subjects play G21 and G21(A) against the *optimizing opponent* and fail to play the dominant solution, they lose the game. Whereas, subjects who play G21 and G21(A) against the *naïve opponent* and fail to play the dominant solution can still win the game. As a result, strategies that do not prescribe playing the dominant solution will always be negatively reinforced by subjects playing the *optimizing opponent.* Therefore, if subjects were using reinforcement learning, then we would expect the decisions of subjects from both the Optimal Player and Optimal Player(A) Treatments to be more in line with the dominant solution, which is largely consistent with what is observed in the data.

Hence, it appears from the data that subjects may be prone to implementing both mimicking and reinforcement learning. Note, however, that in G21 both mechanisms lead to learning the optimal solution when playing the *optimizing opponent*, while in G21(A) only reinforcement leads to learning the optimal solution when playing the *optimizing opponent*. Thus, subjects playing G21 against the *optimizing opponent* can benefit from both learning mechanisms, relative to only one for those playing G21(A). Because of this, we might expect LBD to be relatively more effective in G21, compared to G21(A), when subjects face the *optimizing opponent*. This is largely consistent with the observed data that LBD is only marginally more effective in G21(A) when facing the *optimizing opponent* (refer back to Tables 2 and 5). Similarly, it is reasonable think that player subjects are prone to both mimicking and reinforcement learning, while observer subjects are more prone to mimicking. If this is the case, then LBD should be at least as effective as strategic LBO in G21 when facing an *optimizing opponent*, which is largely consistent with what is observed in the data.

Mimicking and reinforcement might also impact the type of learning that subjects experience when facing the *optimizing opponent*. In particular, mimicking might lead to a so called “eureka” moment where the subject suddenly recognizes the dominant solution of choosing multiples of three. On the other hand, reinforcement learning may be a slower process because subjects discover the winning positions in the game incrementally via some kind of Backward Induction reasoning. Along these lines, we can then think of the Perfect Game measure as a crude proxy for this eureka type learning, and the Total Error Rate measure as a crude proxy for incremental learning. One possible implication that arises is that the observer subjects, who may be relatively more prone to mimicking, may be more prone to eureka type leaning, while player subjects may be more prone to incremental learning. This is consistent with the observed data patterns; namely, Optimal Player subjects have significantly lower Total Error Rates than Naïve Player subjects, and Optimal Observer subjects play significantly more Perfect Games than Naïve Observer subjects.

[23] I refer the reader to Erev and Roth (1998) and Camerer and Ho (1999) for formal models of reinforcement learning.

Please cite this article in press as: Cardella, E., Learning to make better strategic decisions. J. Econ. Behav. Organ. (2012), http://dx.doi.org/10.1016/j.jebo.2012.04.011

4. Conclusion

The motivation of this paper was twofold: The first was to experimentally test whether strategic LBD and strategic LBO are more effective when agents face an opponent who exhibits optimal decision making, relative to sub-optimal decision making. The experimental data reveals 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 only marginally better decision making to those subjects who initially observed the *naïve opponent*.

The second motivation was to experimentally compare the effectiveness of strategic LBD with strategic LBO. The data reveals very little difference in decision making quality between subjects who initially played G21 and subjects who initially observed the play of G21, regardless of the decision making quality of the opponent. These results suggest that strategic LBD is likely no less effective than strategic LBO. By comparison, Merlo and Schotter find experimental evidence that LBO is more effective than LBD in a single-agent profit maximization problem. Because I consider a 2-player strategic game, the player subjects get to "do" *and* they get to "observe" the decision making of the opponent. Hence, there is an element of observational learning, even for agents who play the game. In light of the results of Merlo and Schotter, this might explain why I find that strategic LBD is at least as effective as strategic LBO in the game considered. Investigating which settings are more conducive to LBD and LBO, and the reasons why, remain open questions for future research.

Although the strategic setting considered was abstract and atypical of many strategic settings in the field, a few important insights can be gleaned. First, strategic LBD can be more effective when facing an opponent who makes good decisions, especially in symmetric strategic settings where agents can benefit from both mimicking and reinforcement learning. Thus, agents who are motivated to become better strategic decision makers as quickly and effectively as possible, should consider playing against good decision making opponents. Second, strategic LBD may only be marginally more effective than strategic LBO. This may be particularly relevant in complex settings where repeatedly playing the game may be costly, as it may behoove an agent to gather knowledge by initially observing, in lieu of playing. Hence, under certain settings, apprenticing first and observing good decision making opponents may be as effective at fostering better strategic decision making than actually playing, and in may be potentially much less costly.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.jebo.2012.04.011.

References

Apesteguia, J., Huck, S., Oechssler, J., 2007. Imitation – theory and experimental evidence. Journal of Economic Theory 136, 217–235.
Arrow, K., 1962. The economic implications of learning by doing. The Review of Economic Studies 29, 155–173.
Binmore, K., Shaked, A., Sutton, J., 1985. Testing no cooperative bargaining theory: a preliminary study. American Economic Review 75, 1178–1180.
Binmore, K., McCarthy, J., Ponti, G., Samuelson, L., Shaked, A., 2002. A backward induction experiment. Journal of Economic Theory 104, 48–88.
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.
Camerer, C., Ho, T-H., 1999. Experienced-weighted attraction learning in normal form games. Econometrica 67, 827–874.
Carpenter, J., 2003. Bargaining outcomes as a result of coordinated expectations: an experimental study of sequential bargaining. Journal of Conflict Resolution 47, 119–141.
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.
Dixit, A., 2005. Restoring fun to game theory. Journal of Economic Education 36, 205–219.
Duffy, J., Nagel, R., 1997. On the robustness of behavior in experimental 'beauty contest' games. The Economic Journal 107, 1684–1700.
Dufwenberg, M., Sundaram, R., Butler, D., 2010. Epiphany in the game of 21. Journal of Economic Behavior and Organization 75, 132–143.
Dürsch, P., Kolb, A., Oechssler, J., Schipper, B., 2010. Rage against the machines: how subjects learn to play against computers. Economic Theory 43, 407–430.
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.
Fey, M., McKelvey, R., Palfrey, T., 1996. An experimental study of constant-sum centipede games. International Journal of Game Theory 25, 269–289.
Gneezy, U., Rustichini, A., Vostroknutov, A., 2010. Experience and insight in the race game. Journal of Economic Behavior and Organization 75, 144–155.
Grosskopf, B., Nagel, R., 2008. The two person beauty contest. Games and Economic Behavior 62, 93–99.
Hall, K., 2009. A pure test of backward induction. Doctoral Dissertation, University of Tennessee.
Harrison, G., McCabe, K., 1992. The role of experience for testing bargaining theory in experiments. In: Isaac, M. (Ed.), In: Research in Experimental Economics, vol. 5. JAI Press, Greenwich, pp. 137–169.
Harrison, G., McCabe, K., 1996. Expectations and fairness in a simple bargaining experiment. International Journal of Game Theory 25, 303–327.
Hastie, R., 1986. Review essay: experimental evidence on group accuracy. In: Owen, G., Grofman, B. (Eds.), Information Pooling and Group Decision Making. JAI Press, Westport, CT, pp. 129–157.
Hill, G., 1982. Group versus individual performance: are $N+1$ heads better than one. Psychological Bulletin 91, 517–539.
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.
Huck, S., Normann, H.T., Oechssler, J., 1999. Learning in Cournot oligopoly: an experiment. Economic Journal 109, C80–C95.
Huck, S., Normann, H.T., Oechssler, J., 2000. Does information about competitors' actions increase or decrease competition in experimental oligopoly markets? International Journal of Industrial Organization 18, 39–57.
John, E.R., Chesler, P., Bartlett, F., Victor, I., 1969. Observational learning in cats. Science 166, 901–903.
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.

Jovanovic, B., Nyarko, Y., 1995. The transfer of human capital. Journal of Economic Dynamics and Control 19, 1033–1064.
Laughlin, P., 1980. Social combination processes of cooperative problem-solving groups on verbal intellective tasks. In: Fishbein, M. (Ed.), Progress in Social Psychology. Erlbaum, Hillsdale, NJ, pp. 127–155.
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.
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.
Laughlin, P., Ellis, A., 1986. Demonstrability and social combination processes on mathematical intellective tasks. Journal of Experimental Social Psychology 22, 177–189.
Levine, J., Moreland, R., 1998. Small groups. In: Gilbert, D., Fiske, S., Lindzey, G. (Eds.), In: The Handbook of Social Psychology, vol. 2, 4th edition. McGraw-Hill, New York, pp. 415–469.
Levitt, S., List, J., 2007. What do laboratory experiments measuring social preferences reveal about the real world. Journal of Economic Perspectives 21, 153–174.
Levitt, S., List, J., Sadoff, S., 2011. Checkmate: exploring backward induction among chess players. American Economic Review 101, 975–990.
McKelvey, R., Palfrey, T., 1992. An experimental study of the centipede game. Econometrica 60, 803–836.
Merlo, A., Schotter, A., 2003. Learning by not doing: an experimental investigation of observational learning. Games and Economic Behavior 42, 116–136.
Nadler, J., Thompson, L., Van Boven, L., 2003. Learning negotiation skills: four models of knowledge creation and transfer. Management Science 49, 529–540.
Nagel, R., 1995. Unraveling in guessing games: an experimental study. American Economic Review 85, 1313–1326.
Ochs, J., Roth, A., 1989. An experimental study of sequential bargaining. American Economic Review 79, 355–384.
Offerman, T., Potters, J., Sonnemans, J., 2002. Imitation and belief learning in an oligopoly experiment. Review of Economic Studies 69, 973–997.
Palacios-Huerta, I., Volij, O., 2009. Field centipedes. The American Economic Review 99, 1619–1635.
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.
Rosenthal, R., 1981. Games of perfect information, predatory pricing, and the chain-store paradox. Journal of Economic Theory 25, 92–100.
Schlag, K., 1999. Which one should I imitate? Journal of Math and Economics 31, 493–522.
Shachat, J., Swarthout, T., 2004. Do we detect and exploit mixed strategy play by opponents. Mathematical Methods of Operations Research 59, 359–373.
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. Academic Press, San Diego, pp. 17–45.
Thompson, P., 2010. Learning by Doing. In: Hall, B., Rosenberg, N. (Eds.), In: Handbook of the Economics of Innovation, vol. 1. Elsevier/North-Holland, Amsterdam, pp. 429–476.
Tomasello, M., Davis-Dasilva, M., Camak, L., Bard, K., 1987. Observational learning of tool-use by young chimpanzees. Human Evolution 2, 175–183.
Vega-Redondo, F., 1997. The evolution of Walrasian behavior. Econometrica 65, 375–384.

Please cite this article in press as: Cardella, E., Learning to make better strategic decisions. J. Econ. Behav. Organ. (2012), http://dx.doi.org/10.1016/j.jebo.2012.04.011