A Failure to Communicate:
An Experimental Investigation of the Effects of Advice on Strategic Play*

David C. Cooper
Department of Economics
Florida State University and University of East Anglia
djcooper@fsu.edu

John H. Kagel
Department of Economics
Ohio State University
kagel.4@osu.edu

June 20, 2015

Abstract: We investigate why two person teams beat the truth wins benchmark in signaling games (Cooper and Kagel, 2005 & 2009) by studying an advice treatment where advisees have the benefit of two individuals' insight but not bilateral communication. The TW benchmark states that the performance of a team for problems with a demonstrable correct solution should be no worse than the performance of its most able individual. If one individual solves the problem, the team solves it as well. Advisees whose advisors who play strategically have significantly higher levels of strategic play, but fall well short of the truth wins benchmark as (i) many advisors who play strategically do not provide advice to this effect, and (ii) many advisees fail to follow sound advice. Effect (i) is largely attributable to female advisors who are far less likely to provide advice than men. Whether or not advice contains a clear explanation for its effectiveness has no effect on the likelihood of advisees' strategic play, a result at odds with the idea that economic agents regularly consider all the available alternatives.

Key words: signaling games, strategic play, advice, truth wins, gender effects

JEL classification: C72, C92, D82, L12

* Helpful research assistance has been provided by Nels Christiansen, Kirill Chernomaz, Glenn Dutcher, John Jensenius, David Johnson, Matthew Jones, John Lightle, Peter McGee, Cortney Rodet, Krista Jabs Saral and Xi Qu. Jo Ducey provided valuable editorial assistance. We are thankful for valuable comments received during presentations at Florida State University, Harvard University, University of Maryland, the University of Illinois, Indiana University, the ASSA meetings, the UC Santa Barbara experimental mini-conference, the 2011 SITE conference at Stanford University, and the Society for Judgment and Decision Making meeting in St. Louis. We are grateful for helpful comments from Hal Arkes, Reeshad Dalal, David Budescu, Anthony Buono, David Laibson, Daniel Fragiadakis, and Andrew Schotter. An earlier version of this paper circulated under the title "When are two heads better than one?" This research was partially supported by NSF grants SES-0451981, SES-0924764, and SES-1227298. Any opinions, findings and conclusions or recommendations in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

A number of papers have studied how play of games by teams differs from play by individuals and it is well-established that teams are more strategic than individuals.[1] In earlier papers (Cooper and Kagel, 2005 and 2009) we too report that two person teams in a signaling game have substantially higher levels of strategic play than individuals playing the same game. More importantly, we find that the level of strategic play by teams consistently meets and/or beats a demanding "truth wins" (TW) benchmark (Lorge and Solomon, 1955). The TW benchmark states that the performance of a team for problems with a demonstrable correct solution should be no worse than the performance of its most able individual. If one individual solves the problem, the team solves it as well. Our finding that teams beat this benchmark is surprising since the relevant psychology literature indicates that, while teams typically outperform individuals, it is rare for teams to meet or beat the TW benchmark (Davis, 1992).

The research reported below presents data from new experiments where subjects play the signaling game from our preceding papers in advisor-advisee pairings. Although the TW benchmark is equally valid for advisees as teams, advisees consistently fail to reach this mark. The mechanism underlying this failure helps explain why two person teams beat the TW benchmark in these signaling games. Our results also provide a better understanding of advisor-advisee relationships in strategic environments.

To explore our findings in more detail, we start by explaining the simple model underlying the TW benchmark. There are two distinct branches of the psychology literature on common purpose groups solving problems, one dealing with judgmental questions and another, more relevant for our purposes, dealing with "eureka" problems. Eureka problems require no special expertise or knowledge to solve, but have a demonstrably correct solution that, while not immediately obvious, can be easily confirmed once discovered. Logic problems like the classic Tower of Hanoi puzzle are good examples of eureka problems, and we argue below that strategic play in the signaling game we study also fits well in this category.[2] Consider a team tasked with solving a eureka problem. Suppose it is organized in a way that involves minimal interaction between teammates: team members independently attempt to solve the problem and then

[1] See Charness and Sutter (2012) and Kugler, Kausel, and Kocher (2012) for recent reviews of the literature.

[2] Many commonly studied games (i.e. the beauty contest game, the ultimatum game, coordination games) lack a demonstrably correct course of action and therefore cannot be described as eureka problems. The existing economics literature on team play of games largely studies such games and therefore does not speak to the issue of whether teams meet or exceed the TW benchmark.

compare solutions. As long as the team always recognizes and adopts a correct solution by one of its members (i.e. the truth wins), the probability that the group solves the problem equals the probability that at least one group member solves the problem. Comparing the performance of freely interacting teams with the resulting truth wins (TW) benchmark makes it possible to identify the presence of positive, negative, or zero synergies. In other words, exceeding the TW benchmark implies that working together in a freely interaction team makes subjects better at learning strategic play in the signaling game than they would be working alone.

Our experiments replicate the result that teams beat the TW benchmark and implement two new treatments exploring why. The first is a "self-advice" treatment where subjects play as individuals but were encouraged to type out their thoughts about the game. A central feature of playing in a team is proposing (and possibly explaining) a course of action to your teammate. We hypothesized that the act of writing down their thoughts about the game might stimulate subjects to reason more deeply, leading to the superior performance of teams. If so, self-advice should also lead to more strategic play. The data did not support this conjecture as individuals with and without "self-advice" had essentially the same level of strategic play throughout.

Our second treatment, which is the primary focus of the analysis reported on here, replaces freely interacting teams with advisor-advisee pairings. Advisors and advisees are in fixed pairings and, like teammates, play identical games (i.e. same role and, when relevant, same type). Each member of the pair makes their own decisions and receives the full payoff from their game. Advisors also receive a bonus payment based on their advisees' payoffs as an incentive to provide useful advice. Rather than the bilateral communication available to teams, the discussions of advisor-advisee pairings are limited to one-way communication from advisors to advisees. This permits us to determine whether having the insights of two individuals, which advisees have, increases strategic play sufficiently to meet or beat the TW benchmark. Advisees play strategically more often than individuals, in line with the idea that two heads are better than one, but consistently fail to meet a modified version of the TW benchmark, suggesting that bilateral communication is an essential element of teams' success.

To understand why advisees fall below the TW benchmark, consider the two-step process of giving and receiving advice that underlies the natural extension of the TW model to advisor-advisee pairings. (1) Advisors who play strategically should provide advice to this effect given their strong incentives to help their advisees. (2) Advisees should follow advice to play strategically since they know advisors have incentives to provide useful advice and it is easily confirmed that strategic play is the optimal course of action.[3] In reality, failures occur for both steps of the advice process. For inexperienced sessions, almost half (43%) of advisors who have a history of playing strategically fail to advise their partners to play strategically. This cannot be attributed to a general unwillingness to give advice, as most inexperienced advisors (85%) send messages that include advice about playing the game.[4] Nor can it be attributed to inexperienced subjects who have just considered playing strategically for the first time: 41% of inexperienced advisors who play strategically during the first half of their session *never* advise their partners to do so. The transmission of advice is little improved in experienced subject sessions. While almost all advisors (89%) play strategically in their experienced subject session, and most play strategically relatively early (86% in the first half of the session), 42% of experienced advisors *never* advise their partner to do so. Turning to advisees, a third (34%) of inexperienced advisees who have received advice to play strategically fail to follow it. In short, the TW model assumes that advice acts like a well-made pipe, conveying an advisor's insights to the advisee without any loss, and with advisees acting on this advice. But our data shows that for advisor-advisee pairings this pipe is quite leaky.

In contrast, communication between teammates is both "denser" than between advisors and advisees and less prone to leaks. For inexperienced teams relevant advice was exchanged in 70% of all periods compared to 22% in the advisor-advisee treatment. Further, it is rare that the "truth loses" with teams talking themselves *out* of strategic play.[5] There are a number of conversations between teammates demonstrating that two way communication is necessary to clarify the correctness of the strategic insight passed from one teammate to another.

A key factor behind the failure of advisors who have played strategically to provide salient advice is that women are much less likely to advise strategic play than men, even after controlling for their level of strategic play. The difference is especially striking if we consider the case where advice is most critical, among inexperienced advisors who have played strategically: Seventy three percent (73%) of inexperienced male advisors who have played

[3] In addition, it is assumed that advisees who are not provided with advice are able to figure out and act strategically at the same rate as players acting on their own. This part of the modified TW model is satisfied in the data; i.e., advisees don't stop thinking on their own waiting for insights from their advisors.

[4] Another 8% sent messages that did not consist of advice in an obvious way, and 7% did not send messages.

[5] This occurs 1% of the time.

strategically also give advice to play strategically, compared to only 31% of female advisors. To our knowledge, a gender effect of this sort has not previously been reported in the literature on advice.

With respect to whether advisees follow advice to play strategically, the marginal effect of a sound explanation of why the advice should be followed is essentially zero. Game theorists take it for granted that all available actions will be considered in determining what choices to make, but in a relatively complex game such as ours this assumption is likely to be too strong. In the spirit of rational inattention, individuals playing the game may only consider a small set of seemingly plausible actions (Mussweller et al., 2000; Lord et al., 1984). In this case advice can serve the important role of changing the set of actions under consideration, getting agents to consider the hitherto unrecognized benefits of these alternative actions. This is a surprising result from the standpoint of mainstream economic theory.

Most of the economics literature on advice has focused on how advice can improve performance rather than why it can fail, but there nonetheless exists a large amount of related work in psychology as well as some in economics. Under-responsiveness to advice is commonly observed in the psychology literature where it is referred to as "egocentric advice discounting" (see Bonaccio and Dalal, 2006, for a survey of the psychology literature). However, this literature studies non-strategic environments; e.g., a typical experiment might ask subjects about the dates of various historical events and then provide them with information about other individuals' answers. Low responsiveness to advice has also been reported in interactive settings such as ours in the small economics literature on advice. For example, Schotter and Sopher (2003) study a battle of the sexes game. While they report a surprisingly strong tendency to follow advice, even when it conflicts with the best response to the recipient's beliefs, there remain a substantial minority of subjects who do not follow the advice.[6] What has not been reported in the previous literature is the systematic failure of advisors who themselves have learned to play strategically to offer advice to this effect. We have not seen anything to this effect in the psychology literature. In the economics literature, the closest result comes from Chaudhuri, Schotter and Sopher (2006) who study an individual choice problem, framed as a tournament, with inter-generational advice. About a third of advisors fail to make any specific recommendation to their advisees. Chaudhuri *et al* do not report how well these

[6] In their baseline treatment, 30% of advisees do not follow advice.

advisors performed relative to those who did provide relevant advice, but it seems a safe guess that a number of advisors who chose numbers close to the optimum offered no relevant advice.[7]

In the refereeing process, it was correctly pointed out that the Advice treatment differs from the 2x2 treatment along two dimensions (beyond eliminating bilateral communication) which could affect performance: (1) advisors do not receive the full payoff from their partner's choice and (2) advisors play the game independently from their advisee as well as providing advice. The latter design element plays a central role in generating the main results in our paper, as we would not easily have detected advisors' failures to pass along insights if they had not been playing as well as advising. Nonetheless, it is possible that advisors with a greater stake in their advisee's choice and no distractions due to making their own choice would provide more effective advice, leading to more strategic play by advisees.

We address this possibility by running an "Alternative Advice" treatment where only advisees make decisions, the advisor's sole role is to send messages to the advisee, and the advisor and advisee receive identical payoffs. This led to significantly *lower* levels of strategic play by advisees than in the original advice treatment, as the frequency of advice to play strategically was much lower than in the original advice treatment. The new advice treatment reinforces our conclusion that bilateral communication is necessary for the strong performance of teams in the limit pricing game.

The outline of the paper is as follows: Section I introduces the signaling game which corresponds to a simplified version of Milgrom and Roberts (1982) limit pricing game. Sections II and III summarize the experimental design and procedures. Section IV specifies the main research hypotheses. Section V reports the experimental results both in terms of subjects' behavior in the limit pricing game as well as comparing the content of the dialogues between advisors and advisees and teams. Section VI reports results from the follow-up advice treatment. Section VII summarizes the main results and their potential broader implications.

I. The Limit Pricing Game: The experiment employs a stylized version of Milgrom and Roberts' (1982) entry limit pricing game that focuses on the signaling aspects of the game. An incumbent monopolist (M) faces a potential entrant (E). The game proceeds as follows: (1) M observes its type, high cost (MH) or low cost (ML). It is common knowledge that there are 50%

[7] Chaudhuri *et al* focus on the relationship between receiving advice and advisees' performance, but note in the conclusion that the impact of advice relies on the ability of advisors to communicate their knowledge.

MLs and 50% MHs. (2) M chooses an output (monopolist actions 1 – 7) whose payoff is contingent on E's response. (3) E sees this output, but not M's type, and either enters (IN) or stays out (OUT).[8] Payoffs in experimental currency units are reported in Tables 1a and b.

[Insert Table 1a and b here]

Three features of M's payoffs are worth noting: (1) All else being equal, Ms are always better off if E chooses OUT rather than IN. (2) MLs have a natural tendency to choose higher outputs than MHs. In particular, suppose Ms are myopic, failing to recognize the relationship between their choice and Es likely response. A myopic M will choose 2 as an MH and 4 as an ML. We refer to these choices, highlighted in Table 1a, as the *myopic maxima*. (3) Play of 6 and 7 by MHs (also highlighted Table 1a) are strictly dominated strategies. As such, an ML can perfectly distinguish its type by choosing 6 or 7.

It is always profitable for Es to choose IN versus an MH and OUT versus an ML. Without any information about the M's type beyond the 50-50 prior distribution, it is more profitable for Es to choose IN as shown in the final column (shaded blue).

The asymmetric information, in conjunction with the fact that it is profitable to enter against MHs but not against MLs, provides an incentive for strategic play (limit pricing). The game has two pure strategy separating (sequential) equilibria. In both equilibria MHs play 2 with certainty. MLs play 6 in the efficient equilibrium and 7 in the inefficient equilibrium. In both equilibria Es enter against all outputs that are less than MLs' equilibrium choice. This is supported by out-of-equilibrium beliefs that any output less than 6 in the efficient equilibrium, or less than 7 in the inefficient equilibrium, represents an MH with sufficiently high probability to induce entry. Past research shows that play typically converges, albeit slowly, to the efficient separating equilibrium (Cooper and Kagel, 2005).[9]

Rather than focusing on equilibrium, our data analysis centers on whether MLs play "strategically." *Strategic play is defined as ML's choice of 5, 6 or 7.* While not the only possible definition of strategic play, this is a natural one since these choices are all part of equilibria in which MLs attempt to distinguish themselves from MHs, thereby preventing entry.

[8] We collapse stage 2 of the original Milgrom and Roberts game into the payoffs in Table 1a. This greatly simplifies the experimental design, focusing attention on the signaling aspects of the game.
[9] There is also a mixed strategy equilibrium, with MLs choosing 5 while MHs mix between 2 and 5, which has some drawing power in the data.

Strategic play by MLs fits well with the description of a "eureka" type problem; a problem with a non-obvious but demonstrably correct solution. Playing strategically increases expected payoffs for MLs in all cycles of play for all four treatments. The underlying insight is simple – playing strategically lets MLs prevent entry by distinguishing them from MHs, thereby earning higher payoffs – but is not immediately obvious to subjects. In all of the treatments, the modal choice of MLs in the first cycle of inexperienced play is their myopic maximum (4) even though strategic play yields unambiguously higher expected payoffs. Team dialogues make it clear that this initial lack of strategic play is not a result of mistaken beliefs that choice of the myopic maximum will deter entry, but instead reflects a failure to consider Es' responses to their choices.[10] In other words, subjects initially miss the insight underlying strategic play. Once realized, this insight can be relatively easy to demonstrate to others as the logic is clear and the payoff advantages are overwhelming. Indeed, while advice is not always followed, advisees have significantly higher levels of strategic play than individual play absent any advice.[11]

II. Experimental Design: For all treatments subjects play the limit pricing game for one inexperienced subject session (24 periods) and one experienced subject session (32 periods). Inexperienced subjects have not previously played any version of the limit pricing game. Experienced subjects have played in an inexperienced subject session of the same treatment. The motivation for experienced subject sessions is that it can often take some time for a clean separating equilibrium to develop, as well as to determine if any differences in behavior across treatments are eliminated once subjects gain more experience.

Throughout the paper the term "player" refers to an agent in the limit pricing game, regardless of whether this agent was an individual subject or a pair of subjects playing as a team. The treatments vary who the players are (individuals or teams) and the communication and decision protocol employed.

In the *1x1 treatment*, the players are individuals acting in isolation. They receive feedback about the past history of play, as described below but have no direct input from other subjects.

---

[10] See the detailed coding of team dialogues in Cooper and Kagel (2005) using the same procedures as employed here.

[11] Team dialogues also support the claim that strategic play is self-confirming. When one teammate advocates strategic play while the other does not, invariably strategic play wins, with bad advice winning the days in 1% of all cases where one teammate calls for strategic play and the other argues against it.

In the *2x2 treatment* players are freely interacting two subject teams. Teams are randomly formed at the beginning of a session and remain fixed throughout the session. Teammates must jointly agree on a choice, having (almost) continuous access to a messaging program that allows for bilateral communication about possible actions. Both teammates receive the full payout from their team's outcomes.

In the *self-advice* treatment, play was identical to the 1x1 treatment except for one change: While making choices, subjects were given a text box and prompted to enter "self-advice", period by period, as play evolved. Specifically, the instructions told subjects, "... we have reason to believe that writing out your thoughts regarding what's going on helps in making good choices." Subjects were given continuous access to their self-advice from past plays of the game throughout the session. This treatment was implemented before the advice treatment to see if simply getting subjects to think harder about the game and to verbalize whatever insights they obtained from playing the game, as occurs in the teams' treatment, would be sufficient to replicate the high levels of strategic play found with teams. To foreshadow the results reported on below, there are no significant differences from the 1x1 treatment absent advice, so that we pool the data from the two treatments to provide a baseline against which to evaluate the other treatments.

Subjects in the *advice treatment* were randomly assigned to two subject pairs, with one randomly chosen to be the "advisor" and the other becoming the "advisee." Pairings and roles within a pairing were fixed for the session, as were their roles in any given play of the game. Advisors and advisees played the limit pricing game separately, with no need to agree on a common action (and no mechanism for doing so). Advisors had (almost) continuous access to a messaging program which they could use to send advice to their advisee. The instructions explicitly describe the messages as "advice." Advisees could not communicate with their advisors and were under no obligation to pay attention to the advice provided. Advisors had no way to identify the actions of their advisee, limiting the possibility of non-verbal feedback from advisees. Thus, communication in the advice treatment was strictly unilateral.

To give advisors an incentive to provide useful advice, they received a "bonus" equal to 30% of their advisee's total payoff (along with their own payoff). These bonus payments were only reported at the end of an experimental session so that advisors could not tell what choices their advisees had made. Advisees received their full payoff and had no direct stake in the advisor's outcome.

In any given play of the game an advisor and his advisee played the same role (an M or an E), and were the same type (MH or ML) as Ms. This was common knowledge, with the instructions stressing that an advisor would never play against his advisee. Es were randomly paired with different Ms, with a rotation rule that insured they would not be matched with the same M more than once in each cycle of play.

[Insert Table 2 here]

Table 2 summarizes the number of sessions and number of subjects employed in each of the four treatments, broken down by experience.[12]

III. Experimental Procedures: Inexperienced (experienced) subject sessions used a round robin format where each M met a different E for six (four) plays of the game. Es and Ms then switched roles and played for another set of six (four) games. We refer to the complete set of twelve (eight) rounds as a "cycle." This process repeated itself for a total of 24 (32) games - two cycles in an inexperienced subject session and four in an experienced subject session. The one exception to this rule was experienced subject sessions in the 2x2 treatment. The 2x2 sessions took the longest to run, with play swiftly converging to the efficient separating equilibrium in experienced subject sessions. To avoid keeping subjects for an unusually long amount of time, only three cycles were run instead of four. Given the strong convergence to the efficient separating equilibrium in the second and third cycle, there is no reason to believe that any further changes in behavior would have occurred during a fourth cycle.

At the beginning of inexperienced subject sessions a common set of instructions were read out loud, with each subject having a written copy. This ensured common knowledge of the payoff structure. In particular, advisees knew that advisors were paid a bonus based on their advisee's performance and therefore had strong incentives to provide useful advice. Subjects had copies of both Ms' and Es' payoff tables and were required to fill out a short questionnaire insuring their ability to read them. After reading the instructions, questions were answered out

[12] The 2x2 and advice treatments required that the number of subjects be a multiple of four. In the cases where we were one subject short, one of the lab assistants was employed rather than sending three subjects home and reducing the number of players. In the case of teams the lab assistant announced this fact to her "teammate" along with the fact that they would agree to whatever choice their partner made. For the advice sessions the lab assistant was assigned the role of advisee, and was instructed to follow whatever advice they were given or, in its absence, what the majority of their type was doing. Data for these teams and advisees is excluded from the analysis.

loud and play began with a single practice round followed by a short recapitulation of the instructions as well as another opportunity for questions. At the beginning of experienced subject sessions an abbreviated version of the instructions were read out loud with each subject having a written copy. All instructions were framed in abstract terminology. For example, Ms were referred to as A players, with MHs and MLs described as A1s and A2s. Likewise, Es were called B players and choose between x and y rather than IN and OUT. We use meaningful labels throughout the paper to ease the exposition.

Before each play of the game the computer randomly determined each M's type. This information was prominently displayed on their computer screens. Ms' screens showed payoff tables for both types with the table for their own type always displayed on the left as an additional means of identifying their type. After making their choices the program highlighted Ms' possible payoffs and required that the choice be confirmed. Once *all* Ms had confirmed their choices, each M's choice was sent to the E they were paired with. Es then decided between IN and OUT. Their potential payoffs were highlighted and they were asked to confirm their choices.

Following each play of the game subjects learned their own payoff, the payoff for the player they were paired with, and M's type. In addition, the lower left-hand portion of each player's screen displayed the results of all pairings: M's type, M's output, and E's response ordered by output levels from highest to lowest. The screen automatically displayed the three most recent periods of play, with a scroll bar available to see all past periods.[13] The TW model assumes that subjects can easily confirm that advice to play strategically is correct. Giving them extensive feedback makes it easy to do so since strategic play is overwhelmingly incentive compatible.

1x1 and self-advice sessions employed between 10 - 16 subjects while 2x2 sessions used 16 - 24 subjects resulting in between 8 – 12 players in the team sessions.[14] The number of subjects in the advice sessions varied widely, ranging from a minimum of 12 to a maximum of 28. This variation was intentional. We initially employed relatively large sessions with 20 or more subjects, but this raised a potential confound for comparison with the other treatments as subjects receive feedback from the entire population of players, resulting in up to twice as much feedback

[13] This information was color coded with one color for *all* MLs and another for *all* MHs.

[14] The somewhat smaller team sessions result from the fact that twice as many subjects are needed per player in the 2x2 treatment in conjunction with limited lab space and variation in show-up rates, particularly for experienced subject sessions due to the failure of some subjects to return.

regarding population play as the other treatments. To control for this, we replicated all of the advice sessions using a maximum of 16 subjects. This difference in the number of players had virtually no effect on the results reported and will generally not be dealt with except as a control variable in the regressions.

Subjects were recruited through e-mail announcements directed primarily to introductory economics classes at Ohio State University. Most sessions employed a "double header" design requiring subjects to commit to both an inexperienced and experienced subject session at the same time with the show-up fee ($20), along with half of session one's earnings, withheld until a subject had completed the second session. As part of the process subjects had to agree to come back to one of two scheduled experienced subject sessions, with the inexperienced and experienced sessions conducted within a two week period of each other. This recruiting procedure induced most subjects (87%) to return for an experienced subject session.[15] There are no statistically significant differences in the likelihood of strategic play by subjects who did not return and those who did. Since subjects could choose between two experienced subject sessions, these sessions contained a mix of subjects from different inexperienced sessions. By the same token, it was not possible to keep pairings or roles fixed across sessions.

Inexperienced (experienced) subject sessions lasted a little under two (one and a half) hours. All payments were in cash, with the experimental currency ("francs") converted to dollars at a 400:1 ration. Earnings averaged slightly less than $69 for a subject who completed both sessions, including the $20 show-up fee.

IV. Initial Hypotheses: The results section focuses on comparing levels of strategic play of MLs as advisors and advisees, as well as comparing MLs' strategic play in the advice treatment with the 1x1 and 2x2 sessions. This section formulates initial hypotheses about these treatment effects by extending the TW model to the advisor-advisee sessions. Let p(T) be the probability that an individual in treatment T $\in$ {1x1, 2x2, self-advice, advisor, advisee} plays strategically as an ML *independent from being told the solution by another subject*. The probability of an individual with access to input from a group of size N playing strategically in treatment T is 1 -

[15] As a methodological check, we ran half of the 1x1 sessions using standard recruiting methods with no special inducement to return. Not surprisingly, the percentage of subjects returning was substantially lower (65%) than in the "double headers" (95% in the 1x1 treatment, 87% overall). The data shows that the recruiting method did not systematically affect the likelihood of MLs playing strategically.

$(1 - p(T)*(1-PL))^N$, where $PL \geq 0$ represents "process loss". If PL = 0 this is the standard TW benchmark where a group solves the problem if any individual solves it.

The process loss term, PL, has a very specific interpretation in this generalized version of the TW model. It does *not* represent changes in individual behavior due to treatment effects – these are captured through the p(T) function. Thus, factors that are natural examples of process loss in teams, such as distraction due to irrelevant conversation or reduced effort from free-riding, are incorporated into p(T) rather than PL. Instead, PL relates directly to the assumption underlying the TW model, that a group solves the problem if any of its members solves it. Even with a problem where the solution is verifiable, this assumption might not hold as individuals may fail to communicate a solution to others or fail to persuade others that the solution is correct. *PL gives the probability that an individual who has successfully solved the problem fails to successfully communicate the solution to the other group members.* As noted previously, this rarely happens in the 2x2 treatment. In formulating hypotheses below, we adopt the simplifying assuming that PL = 0 for both the 2x2 and the advice treatment, but in discussing the results of the advice treatment we relate performance by advisees to strong evidence that PL > 0.

The probability of playing strategically as an ML in the 1x1 treatment reduces to p(1x1). If p(2x2) = p(1x1), the probability of an N person team solving the problem reduces to the standard TW benchmark: $1-\left(1-p(1x1)\right)^N$. If teams meet or beat the TW benchmark it follows that $p(2x2) \geq p(1x1)$ in spite of the factors (free-riding, distraction) mentioned above.

The self-advice treatment does not require subjects to convince a teammate to follow a particular course of action. The probability of solving the problem therefore reduces to p (self-advice). We are unaware of any results in the psychology literature, one way or the other, on this point.[16] As such we adopt a null hypothesis of no effect on behavior as a result of self-advice.

Self-Advice H1: Self-advice has no effect on the probability of strategic play compared to a standard 1x1 treatment in which players had no opportunity to comment on their strategies and thoughts: p(self-advice) = p(1x1).

For the advice treatment we distinguish between advisors and advisees. The null hypothesis for advisors is:

[16] "Talk aloud" protocols are the closest analogue we are aware of to our self-advice treatment (Ericsson and Simon, 1980).

Advisor H1: Giving advice has no effect on the probability of strategic play: p(advisor) = p(1x1).[17] The frequency of MLs' strategic play as advisors equals the frequency of strategic play in the 1x1 treatment, with or without self-advice.

H1 is based on the idea that since advisors are essentially working on their own (only one head to work with) there is no reason to believe that they will play any more strategically than subjects in either the 1x1 or the self-advice treatments. There are several reasons why this null hypothesis might not hold.[18] First, the psychology literature on accountability suggests that self-critical and effortful thinking is most likely to be activated when decision makers know that they will be accountable to an audience that is interested in accuracy and is reasonably well informed (Lerner and Tetlock, 1999), as advisors were. Other factors that might promote higher levels of strategic play include (i) increased monetary incentives as advisors earnings include a modest percentage of their advisees earnings (however, the literature suggests that such a modest increase in incentives will have no major impact on outcomes; see Smith and Walker, 1993, for example) and (ii) to the extent that advisees, on account of having two heads to work with, play more strategically as Ms and are more sensitive to the incentives to enter as Es (increasing the entry rate differential between 4 and 6), advisors receive more informative feedback and face stronger incentives to play strategically than subjects in the 1x1 and self-advice sessions.

We propose two alternative hypotheses for advisees.

Advisee H1: Receiving advice has no effect on the probability of strategic play: π(advisee) = π(1x1) = p(1x1).

Advisee H2: An advisee plays strategically if his advisor plays strategically; otherwise he plays strategically with the same probability as a subject in the 1x1 treatment:

$\pi(\text{advisee}) = 1 - (1 - p(\text{advisor}))*(1 - p(1x1)) > p(1x1)$.

Advisee H1 is the null and is based on the idea that advisees are under no obligation to follow the advice they receive and can choose to completely ignore it. If so, they should play no differently than subjects in the 1x1 treatment. Advisee H2 draws on the logic of the TW model. An advisor who figures out how to play strategically as an ML should recognize the value of this

[17] Alternatively p (advisor) = p (self-advice). To foreshadow the experimental results and to simplify the specifications, in what follows we drop the distinction between p (1x1) and p (self-advice).

[18] Greater strategic play by advisors would be consistent with the results of Iyengar and Schotter (2008), who report that in a decision theoretic context advisors, who did not make choices themselves but were paid on the basis of their advisees' choices, came closer to the optimal choice than subjects deciding alone (either as advisors or advisees).

insight and pass it on to their advisee. Advisees who receive advice to play strategically should follow it. In other words, there should be no process loss. An advisee might also figure things out on their own absent any advice to play strategically. Assuming that receiving bad (or at least not useful) advice has no effect on their behavior, the probability of an advisee figuring out strategic play on their own should be the same as in the 1x1 treatment. Taken together, these generate the version of the TW benchmark given in Advisee H2. Advisee H2 predicts *higher* levels of strategic play for advisees than in the 1x1 treatment. This does not reflect advisees being any better at figuring out how to play strategically, but instead is strictly a function of them having two heads (their own and their advisor's) to rely upon.

V. Results: The efficient separating equilibrium gradually emerges in all four treatments. Figure 1 shows data from the 1x1 treatment, but the dynamics are similar across treatments, varying primarily in speed of convergence. MLs start out largely ignoring the strategic possibilities of the game, choosing their myopic maximum (4), even though entry rates make strategic play (MLs choosing 5, 6 or 7) incentive compatible by a wide margin from the start. Over time MLs learn to distinguish themselves from MHs by choosing 5 or, more frequently, 6. Our focus is on the development of strategic play under the different treatments.

[Insert Figure 1 here]

5.1 *1x1 versus the self-advice treatment:* Figure 2 compares strategic play for MLs in the 1x1 and self-advice treatments. The differences between the two treatments are modest, and the regression shown in Table 3 reports no statistically significant differences. This regression is a logit model with nested random effects at the session and team/individual level. The dataset includes all ML observations from the 1x1 and self-advice treatments. The dependent variable is a dummy for strategic play and the independent variables include dummies for the current cycle, interactions between the cycle dummies and a dummy for the self-advice treatment, the number of subjects in the session, and the entry rate differential. The latter is the difference in entry rates between 4 and 6 for the current cycle. It is calculated across all observations from the same session for the current cycle, since subjects see all entry decisions in their session. The entry rate differential controls for the incentives to play strategically.

The variables of interest are the interactions between the cycle dummies and the self-advice treatment. These capture the difference in each cycle between the 1x1 and self-advice treatments. None of these six variables is statistically significant, nor are the six variables jointly significant (χ^2 =8.49; p = 0.205). Given this lack of statistical significance, in what follows we pool the data from the 1x1 and self-advice treatments for MLs.[19]

[Insert Figure 2 and Table 3 here]

A natural guess for why self-advice has no effect is that subjects might have written little self-advice since there was no direct incentive to do so and subjects had only our word that it would be helpful. In fact, use of self-advice was common for inexperienced subjects as 65% entered at least one piece of *relevant* self-advice (how to choose as either an M or E, or in reference to the population data provided), with relevant self-advice being recorded on average in 4.8 rounds of inexperienced subject play.

Conclusion 1: MLs strategic play in the self-advice sessions is statistically indistinguishable from the level of strategic play in the 1x1 sessions.

[Insert Figure 3 and Table 4 Here]

5.2 *1x1 versus the 2x2 treatment*: Figure 3 compares MLs' strategic play in the 1x1 and 2x2 treatments. The TW benchmark is also shown, including error bars for the 90% confidence interval.[20] MLs' strategic play is higher in the 2x2 treatment than in the 1x1 treatment for all cycles, with teams achieving 100% strategic play in the second and third cycles of experienced subject play. Teams meet or beat the TW benchmark in each cycle of play, replicating earlier results for the same game (Cooper and Kagel, 2005). As noted earlier, the fact that teams act more strategically than individuals is a common finding (see Charness and Sutter, 2012, and Kugler, Kausel, and Kocher, 2012, for recent reviews of the literature), but it is far less common for them to meet or beat the TW benchmark (Davis, 1992).[21]

---

[19] Our conclusions are little affected by alternative specifications, such as use of a linear probabilities model or estimating the treatment effects without controls for the session size and entry rate differential.

[20] The formula for the truth wins bench mark based on play in the 1x1 treatment is $1 – (1 – p(1x1))^2$. Because of clustering in the data, simulations are needed to correctly calculate the error bars. The simulated 2x2 data is based on 250,000 simulated 2x2 data sets for each cycle of play, with the same number of teams in each data set as in the experiment. Simulated 2x2 play is based on randomly drawing two subjects (with replacement) from the 1x1 sessions. A simulated team was considered to have played strategically if either of its members played strategically. The error bars then display the 5th and 95th percentiles of the distribution of percentages of strategic play in a simulated 2x2 data set.

[21] The level of strategic play here is substantially higher for *both* the 1x1 and the 2x2 treatments than reported in Cooper and Kagel (2005), even though we are using the same procedures and the same subject pool. The proximate cause for this appears to be substantial increases in SAT/ACT scores and high school class rank of entering freshmen at Ohio State. We plan to explore the relationship between cognitive ability and strategic play in another paper (Cooper and Kagel, work in progress).

The regression reported in Table 4 documents that the difference between the 1x1 and 2x2 treatments is statistically significant. The dataset includes all ML observations from all treatments. The dependent variable is strategic play (choice of 5, 6, or 7). Independent variables include dummies for the current cycle, interactions between the cycle dummies and treatment dummies, the number of players in the session, and the entry rate differential between 4 and 6. The 1x1 and self-advice data have been pooled and will be referred to as the "1x1 data," which serves as the baseline against which to compare the 2x2, advisors, and advisees. The variables of primary interest are interactions between the cycle dummies and the treatment dummies. These capture the difference between the treatments and the 1x1 data. The number of players reflects the number of individuals for all treatments except the 2x2 treatment where it reflects the number of teams – the number of players in the advice treatment was systematically varied to determine whether differences between the advice and 2x2 treatments are due to differing amounts of feedback. The entry rate differential between 4 and 6 is calculated in the same fashion as described for Table 3. Random effects are at the session and individual/team level.[22] Focusing on the interaction terms for the 2x2 treatment, all are significant at the 5% level or better.

Conclusion 2: MLs' strategic play in the 2x2 sessions is consistently higher than in 1x1 sessions and meets or exceeds the TW benchmark in each cycle of play. This replicates the results reported in Cooper and Kagel (2005).

[Insert Figure 4 Here]

5.3 *1x1 versus the advice treatment:* Figure 4 compares MLs strategic play in the 1x1 and advice treatments. Data for advisors and advisees is reported separately. Advisors' initial level of strategic play is slightly higher than in the 1x1 treatment, and the gap increases for the second cycle of inexperienced play. This difference is eliminated for experienced subjects, with advisors' strategic play slightly *lower* than in the 1x1 treatment in the last three cycles of experienced subject play. Advisees consistently do better than advisors with the difference varying little over time. Advisee's strategic play is higher than in 1x1 games for both inexperienced cycles, especially Cycle 2, but dips slightly below the 1x1 treatment in the last two cycles of experienced play.

[22] For the advice treatment, the second random effect is at the level of an advisor-advisee pair rather than the individual advisor or advisee. Even if no advice is sent, advisors and advisees share a common history of roles and types which may induce correlation, so we take a conservative approach. Once again, the results are robust to the use of alternative specifications.

Returning to the regression analysis reported in Table 4, strategic play by advisors is only significantly higher than in the 1x1 treatment for the second cycle of inexperienced play (and is significantly lower for the final cycle of experienced play). For advisees there is significantly more strategic play than in the 1x1 treatment for the second cycle of inexperienced play ($p < 0.01$) and (weakly) for the first cycle of experienced play ($p < 0.10$). At no point does strategic play for advisees dip significantly below levels for the 1x1 treatment.[23]

Advisees' strategic play relative to *their* TW benchmark (Advisee H2) is also shown in Figure 4 along with error bars for the 90% confidence interval. To reiterate the basic idea, an advisee should play strategically if either his advisor plays strategically (the TW model assumes that an advisor always passes on advice to play strategically and this advice is always followed) or if an advisee figures it out on his own (which is assumed to happen with the same probability as in the 1x1 sessions). The TW benchmark is *not* met as advisees' strategic play is always below it, with the differences statistically significant in all but the second cycle of inexperienced subject play.

Conclusion 3: Both advisors and advisees have higher levels of strategic play than in the 1x1 treatment through the first cycle of experienced subject play, at which point strategic play tapers off with the 1x1s catching up. Advisees play strategically more often than advisors in all cycles of play, but less than their (modified) version of the TW benchmark.

Figure 4 also compares strategic play in the 2x2 games with advisors and advisees. For all cycles, strategic play is greater in the 2x2 treatment than for either advisors or advisees. This difference is significant for advisors in both inexperienced (at the 10% level) and experienced subject play, and for advisees in experienced subject play.

[Insert Figure 5 Here]

5.4 *Process Loss in the Advisor-Advisee Treatment:* The failure of advisees to meet the TW benchmark is illustrated by Figure 5. Their TW model assumes that (1) an advisee matched with an advisor who does *not* play strategically should play strategically no less than subjects in the 1x1 treatment (since they might figure out strategic play on their own), and (2) advisees whose advisors play strategically should always play strategically since their advisor communicates the

[23] The difference between advisors and advisees is significant in the second cycle of inexperienced play, as well as the second and third cycles of experienced play.

relevant information and the advisee follow the advice. Part (1) holds for inexperienced subjects, as there is no significant difference in strategic play between the 1x1 treatment and advisees whose advisors are not playing strategically. However, part (2) fails for inexperienced subjects, as slightly more than a quarter of the advisees whose advisors played strategically did not play strategically themselves. While these advisees play strategically significantly more often than in the 1x1 treatment as well as those advisees whose advisors did *not* play strategically, they do not come close to the 100% strategic play that Part (2) of their TW model requires. Part (2) of the TW model also fails for experienced subjects: Advisees whose advisors play strategically play strategically 90% of the time. However, given the high rate of strategic play for experienced subjects (83%) in the 1x1 treatment, this 10% failure rate is sufficient for the TW benchmark not to be met.

To better understand the failure of the advisor-advisee pairings to meet the TW benchmark, we coded the content of the advice given in relationship to subjects' behavior. Advisors' messages were coded using two simple categories: (1) Was a subject advised to play strategically as an ML (i.e., given explicit advice to choose 5, 6 or 7); and (2) Was a subject given an explanation of why he/she should play strategically as an ML. An explanation of strategic play was broadly defined – it was sufficient, for example, to say that strategic play would make entry less likely without providing any reason this might be true. Two research assistants scored the first time (if ever) a comment in each category was made. Agreement between the two coders was high – they agreed on the first occurrence in Categories 1 and 2 for 85% and 90% of all observations respectively. We also coded how many times each advisor provided relevant advice, where relevant advice is defined as any message on how to play as an M or E.[24] Before discussing the results of this coding exercise, it is useful to clarify some terminology: If we say that an advisee has "received advice to play strategically in Period t," this means that the first time they were advised to play strategically as an ML occurred in Period $s \leq t$ of that session. In other words, once an advisee has received advice to play strategically, they are assumed to retain that advice for the remainder of the session. This is a reasonable assumption as advisees can access past messages, advice to play strategically is never contradicted in later periods, and the nature of advice makes it unlikely that advisors will feel the need to repeatedly give the same

[24] Given the straight forward nature of this task only a single coder was used.

advice. Extending this terminology, when we state that xx% of advisees in Block XX received advice to play strategically, this means that xx% of the observations of advisees in Block XX received advice to play strategically. Further, stating that their advisor has played strategically in Period t means that the advisor *first* played strategically as an ML in Period s ≤ t.

5.5 *Process Loss Attributable to Advisors:* Most advisees received at least some relevant advice – 85.4% (78.4%) of inexperienced (experienced) advisees received at least one piece of relevant advice during an experimental session. Relevant advice was fairly common in general, with inexperienced (experienced) subjects averaging 7.3 (7.4) relevant messages per session. However, even though 70% of inexperienced advisors played strategically at least once (chose 5, 6 or 7), in only 57% of these cases was the advisee explicitly told to play strategically.

A number of potential explanations for this breakdown in advice giving have been suggested. One possibility is that advisors who fail to give advice have less experience with strategic play than those who give advice. However this is not the case: Among inexperienced advisors who play strategically at least once, those who never send advice to play strategically are themselves playing strategically an average of 3.6 times per session, only slightly less than the average of 3.8 times for those who do send advice to play strategically.[25] As another way of making this point, consider advisors who first play strategically in the first half of the session (Cycle 1 for inexperienced sessions, Cycles 1 and 2 for experienced sessions). In inexperienced (experienced) sessions, 41% (38%) of these advisors *never* advise their partners to play strategically.[26] As such the failure to give useful advice reflects something deeper than a lack of understanding or opportunity.

Second, it's been suggested that by giving advisors a bonus based on their advisee's earnings, we may have inadvertently incentivized advisors to hedge their payoffs, playing strategically while advising otherwise for their advisee. We checked for this explicitly in each of the nine inexperienced subject sessions. We could not find a *single* instance in which an advisor who had played strategically gave advice to play non-strategically.[27] This lack of hedging makes sense since the incentives for MLs to clearly distinguish themselves from MHs are so strong from very

[25] The corresponding numbers for experienced subjects are 6.6 and 7.8 strategic plays per session.

[26] In Cycle 2 of inexperienced play, these advisors played strategically 81% of the time as MLs. The equivalent figure for Cycles 3 and 4 of experienced play is 84%.

[27] We also checked if any advisors who chose 5 advised choosing 6, or vice versa. This never happened.

early on that advisors would have to be quite risk averse to hedge in this way. There are also obvious social conventions which likely inhibit advisors from hedging.

Third, it was suggested that part of the explanation for this breakdown in advice giving could be gender based. Given that women are less confident in their abilities than men in a number of strategic settings (e.g. Niederle and Vesterlund, 2007), it seemed plausible that this lack of confidence would translate into unwillingness to give advice. Much to our surprise this turns out to be the case as men are much more likely to give advice to play strategically than women in all cycles of play, both unconditionally and conditional on having played strategically themselves. In the first cycle of inexperienced subject play, 72% of the men who played strategically also gave advice to this effect, whereas *none* of the women who played strategically did so. This gap narrows over time, but even in experienced subject sessions, women who have played strategically never get above 50% for offering advice while men are consistently at 70% or higher.

[Insert Table 5 here]

Table 5 presents results from Cox hazard rate models formally confirming that men and women differ in their speed to offer advice. The dataset includes all advisors whose gender is known. The dependent variable is the first time advice is given to play strategically (censored if advice hasn't been given when the session ends), resulting in one observation per advisor per session. We report the coefficient values and standard errors (in parentheses) for each independent variable. A positive (negative) coefficient implies that the advisor will be faster (slower), ceteris paribus, to provide advice to play strategically. Since an advisor can appear in the dataset twice (as inexperienced and experienced), the standard errors are corrected for clustering at the advisor level. All three models include controls for the advisor's gender, whether the advisor is experienced, the number of advisors in the session (session sizes were systematically varied), and the entry rate differential.

Model 1 is a basic regression with no further controls. The gender dummy has a negative sign, which is significant at the 5% level.

Model 2 adds a control for advisors' experience with strategic play (the percentage of times an advisor chose to play strategically as an ML, *scaled from 0 to 1*). By controlling for differences in the frequency of strategic play, we rule out the obvious explanation that gender *differences* result from women being less likely to play strategically than men in the signaling game. Not surprisingly, advisors who play more strategically are also significantly faster to offer advice to play strategically. The gender dummy remains statistically significant, albeit at the 10% level.

Model 3 adds an interaction effect between the percentage of strategic play and gender. The gender dummy is now positive, while the interaction between strategic play and gender is negative. Both gender variables, the dummy and interaction terms, are significant at the 1% level, and are jointly significant at the 1% level ($\chi^2 = 13.1$). Looking at the two gender variables together, as the proportion of strategic play approaches 100%, Model 3 shows that women are *less* likely to provide advice to play strategically $p < 0.01$), consistent with the raw data in.[28] Interestingly, the model also shows that with 0% strategic play women are *more* likely to provide advice to play strategically than men (the significant positive coefficient for the gender dummy). This too is consistent with the raw data: inexperienced (experienced) women advisors who have not played strategically provide advice to this effect 16% (17%) of the time versus 5% (2%) for men. In short, comparing advisors who have played strategically, men are more likely to provide advice to this effect than women. But comparing advisors who have not played strategically, women are more likely to provide correct advice to play strategically than men. Both effects point to men being more confident in their beliefs regarding strategic play than women.

As a final piece of evidence, there are no significant or consistent differences between men and women in the extent to which they provide bad advice (advice to *not* play strategically as an ML), ruling out the notion that men are simply more prone to providing advice, good or bad, as an explanation for the gender effect reported.

Conclusion 5: A major cause of advisees' failure to meet TW benchmark is that advisors who play strategically fail to give advice to this effect. Underlying this failure are significant differences between men and women in the frequency of providing advice to play strategically. Women are significantly ***less*** *likely to provide such advice conditional on having played strategically.*

Our findings about advisors' failures to provide useful advice, as well as the strong gender differences in the provision of advice, have not been previously identified in either the

[28] Formally, we redo Model 3 with the ratio for strategic play set to 1 in place of % Strategic Play. The gender dummy now captures the difference between women and men at 100% strategic play. The coefficient is negative (-1.102) and significant at the 1% level (standard error = .370).

economics or psychology literatures. We conjecture that the difference between men and women is driven by men's greater confidence in their insightfulness. Given that there is no doubt substantial variation in confidence levels within our sample of men and women, this implies that a lack of confidence accounts for the failure to advise strategic play after having played strategically regardless of gender.

5.6 *Process Loss Attributable to Advisees:* In about a third of all cases advisees who were given sound advice to play strategically failed to follow that advice. As noted earlier this result parallels findings from the psychology and economics literature on advice taking and decision making where one of the most robust findings is that advisees do not follow advisor's recommendations nearly as much as is optimal, referred to in the psychology literature as egocentric advice discounting. The failure of advisees to follow good advice is especially puzzling in our experiment since advisees knew that advisors had financial incentives to provide useful advice and received extensive population feedback that allowed them to easily verify that strategic play was optimal.

A surprising element of advisees' behavior, with no existing parallel in the literature on advice, is that being given an explanation as to why they should play strategically has essentially no more impact on inexperienced advisees than simply being told to choose 5, 6, or 7. Table 6 documents this for both inexperienced (top panel) and experienced subjects (bottom panel). Observations are distinguished by whether the advisee had *ever* been told to play strategically and whether the advisee had *ever* been given an explanation for playing strategically.[29] For each cell we report the frequency of strategic play along with the number of observations in each cell (in parentheses).

[Insert Table 6 here]

For inexperienced subjects there clearly is no differential effect from being provided with an explanation. Although it appears that providing an explanation has a positive marginal impact for experienced subjects, the econometric analysis reported below shows that the marginal impact of an explanation is not statistically significant for either inexperienced or experienced advisees. The fact that there is no significant increase in the likelihood of strategic play after

[29] While rare, there are a few instances in which advisees' were given an explanation for why they should play strategically without ever being given explicit advice to choose 5, 6 or 7. For example, one advisor wrote, "You want to make the [entrant] think that you are [an ML] so they will choose [Out]." However, this advisor failed to explicitly tell the advisee to choose 5, 6 or 7 to achieve this outcome.

being given an explanation suggests that the primary effect of advice is to make subjects consider strategic play.[30] Economists and game theorists typically assume that decision makers consider all possible actions, but the psychology literature indicates that they routinely neglect some possibilities. External prompting can get decision makers to actively consider otherwise ignored options (see, for example, Mussweller et al., 2000 and Lord et al., 1984).[31] In our experiment, entry is much higher on 2 than on 4 in the first several plays of the game. This tends to reinforce MLs' initial beliefs that their myopic strategy (4) is best. As a result they become locked in on 4 and fail to notice that choice of 6 is actually a better option. As such the primary role of advice appears to be to get MLs to reconsider an option that they had previously dismissed.

The low impact of receiving an explanation makes the occurrence of egocentric advice discounting especially striking in our experiment. The entry rate differential on 4 versus 6 is such that there is a demonstrably correct course of action in the limit pricing game. Yet, even with the benefit of an explanation, about a third of inexperienced advisees fail to follow advice to play strategically.[32]

[Insert Table 7 here]

Table 7 provides formal backing for these assertions. The dataset is limited to observations from advisees playing as MLs and the dependent variable is a dummy for strategic play. The independent variables of primary interest are dummies for what sort of advice has been received, either in the current period or previously in the session. We account for three types of advice: (1) has the advisee been advised to play strategically as an ML, (2) has the advisee received an explanation for why he/she should play strategically as an ML, and (3) has the advisee received

[30]Another possible reason that explanations fail to have a significant positive effect is that many explanations are poor quality. We did not have coders rank the quality of explanations, judging this to be excessively subjective, but it was obvious that the quality of explanations ranged widely. However, to claim that the lack of an effect is based strictly on the low quality of explanations is implausible. The many cases where the explanation was minimal (e.g., "I think we pick 6 as a2s [MLs] and 4 as a1s [MHs] that way we insure they pick y [out] when we're an a2") are balanced by the many cases where the explanations provided were quite good ("obviously if they pick a 6 or 7 they are an A2 [ML] because no one will pick a negative number…..if you're an A2 pick 6 because they will know you're an A2 and no matter pick y [out]).

[31] Schotter (2003) makes a similar conjecture based on data from a numbers game against nature with an observer providing written advice to the decision maker. Our data strengthens the evidence for this conjecture by showing that an explanation for use of strategic play is not necessary for advice to have a positive effect on outcomes.

[32] On average, advisees who received advice to play strategically earned 15 francs more per period as an ML than those who had not received such advice. The effect of good advice on advisees' payoffs is limited by their failure to follow it.

bad advice. Bad advice refers to cases where the advisee has specifically been told not to play strategically as an ML (i.e., choose 4) and has not received later advice countermanding the suggestion to play non-strategically. The three dummies for types of advice are interacted with dummies for inexperienced andexperienced play. The regression also includes dummies for the current cycle of play. Random effects are at the session and individual/pairing level.

For inexperienced play, there is a positive and significant effect on strategic play from having received advice to play strategically and a negative and (weakly) significant effect from bad advice. The negative effect of bad advice makes it even more difficult for advisees to reach the TW benchmark since the model assumes that incorrect advice is ignored. The effect of receiving an explanation for advice to play strategically is positive, but small and nowhere close to statistical significance. The results for experienced play are similar but weaker. Advice to play strategically again has a significant effect, but only at the 10% level. Oddly, bad advice has a positive effect on strategic play, albeit nowhere close to statistical significance. Once again there is little effect from receiving an explanation for why the advisee should play strategically.

Conclusion 6: The marginal impact of providing an explanation as to the logic underlying why their advice should be followed has no appreciable impact on the likelihood that the advice will be followed. As such the primary impact of advice appears to be to get subjects to consider alternative strategies that they have come to ignore.

Having documented a strong gender effect in the giving of advice, we looked into whether there might be strong gender differences in acting on advice to play strategically. The effect of advice to play strategically is consistently larger for woman than men. However, these differences between the genders do not achieve statistical significance in the data.[33]

5.7 *Process with Advice versus 2x2 Teams:* Consistent with the superior performance of teams in the 2x2 treatment, there are substantial differences in the amount of relevant information exchanged in the team discussions compared to the monologues in the advice treatment. In the inexperienced 2x2 sessions, relevant advice was exchanged in 70% of all periods compared to 22% of all periods in the inexperienced advice treatment.[34] The corresponding percentages for

[33] We have modified the model in Table WW by adding dummies for inexperienced and experienced women as well as interactions between these dummies and dummies for having received advice to play strategically. The interaction terms are both positive but are not statistically significant for either inexperienced (z = 0.79) or experienced (z = 1.47) women.

[34] All of the figures in this paragraph come from the limited coding of the 2x2 data reported noted previously, which was done by a single coder.

experienced subjects, although lower, are also quite different: 34% of all periods in the 2x2 sessions versus 16% in the advice sessions. All of the inexperienced teams exchanged at least one piece of relevant information at some point during an experimental session compared to 85% for inexperienced advisors.[35] Most importantly, 46% of inexperienced advisors explicitly suggested play of 5, 6, or 7 compared to 78% of inexperienced teams discussing the benefits of such choices. The substantially more complete information exchange in the 2x2 treatment no doubt contributes to the fact that teams meet or beat the TW benchmark compared to advisees who consistently fail to meet their (modified) TW benchmark.

This free flowing two way communication also serves to clarify the correctness of a strategic insight from one teammate to another, which is not available in the advisor-advisee treatment. The following discussion illustrates one of a number of such cases we have identified, in this case for an ML team deciding to choose 6:

10: "do you want to go with 6?"
14: "what about 4? Okay"
10: "if we choose 4 they will choose x for sure"
10: "if we go with 6 we are guaranteed 592"
14: "why? they can choose y (OUT) for 50 more"
14: "and we both get bigger one"
10: "what do you mean?"
14: "if we choose 4, they will know we are a2 because 4 is the biggest in a2" "and in a2 y is bigger for them"
10: "but everytime 4 has been chosen, each time the b response has been x (IN)"

Teams resemble the advice treatment along one important dimension, gender. When subjects play as individuals, men are slightly more likely than women to play strategically as MLs.[36] This effect is much more dramatic in the 2x2 treatment. For inexperienced subject session, teams with two men play strategically almost twice as frequently as those containing at least one woman (84% vs 46%), with this difference statistically significant.[37] Teams with two men are also (weakly) significantly faster to discuss strategic play.[38] Given the small difference between

[35] Corresponding values for experienced subjects are 90% for teams versus 78% of advisors.
[36] Across treatments where subject choose as individuals (1x1, advisors, advisees) we observe more strategic play by men than women. The difference isn't large, roughly 10%, but consistent and, at times, statistically significant.
[37] This result is based on a relatively small sample of 31 teams, 12 of which were all male. ,
[38] On both strategic play and speed of discussing strategic play, there is little difference between teams with one or two women. These differences disappear with experienced subjects. Since virtually all teams play strategically when experienced, there is little room for an effect to occur.

men and women playing as individuals, we suspect that the gender effect within teams has similar roots as the gender effect with advisors. That is, we conjecture that men are more confident about their insights than women and are therefore more willing to share them within a team, without which one loses much of the "two heads" effect.

Conclusion 7: Strategic play is greater in the 2x2 treatment than for either advisors or advisees. There are substantially higher levels of relevant communication between teammates than from advisors to advisees which, among other things, serves to reinforce the correctness of the strategic insights from one teammate to another.

VI. Alternative Advice Treatment: One important purpose of our experiment was to explore why teams are able to beat the TW norm in the limit pricing game. The poor performance of advisees relative to the TW norm implies that bilateral communication is a necessary element for teams' strong performance, but it is possible that other elements of the experimental design have artificially reduced the performance of advisees. The advice treatment differs from the 2x2 treatment along two dimensions which could affect performance: (1) advisors do not receive the full payoff from their partner's choice and (2) they are playing the game independently from their advisee as well as providing advice. There were important reasons for these design choices. If advisors had not been playing as well as advising, we would not easily have detected their failure to pass along insights to play strategically as an ML. Nonetheless, it is possible that advisors with a greater stake in their advisee's choice, and no distractions due to playing their own game, would provide more effective advice resulting in more strategic play by advisees. We address this possibility by running a follow-up treatment, the "Alternative Advice" treatment.

The Alternative Advice treatment is a hybrid of the Advice and 2x2 treatment. As in the 2x2 treatments, subjects are formed into fixed two person teams that play a single game in each round. Each of the partners receives the full payoff from this game. As in the Advice treatment, teammates are assigned fixed roles as either an advisor or an advisee. At the beginning of each round, the advisors are provided with a text box where they can type a message to their advisee. The advisees see these messages prior to choosing an action, but cannot respond. The advisors have exactly the same information as the advisees. They know whether they are an ML, MH, or E, and as an E they know the choice made by the M player. The decision for the team is made by the advisee. Since advisors' payoffs are dependent solely on the actions of their advisees, they presumably have strong incentive to give good advice. At the end of each round, both advisors and advisees receive the same feedback as in the Advice treatment. They both see the full set of outcomes for all pairs in a period as well as a history window showing previous outcomes. For advisors, the outcome for their advisee is not specifically identified (although it often can be inferred). As in the Advice treatment this was done to cut a possible back channel of communication from advisees to advisors.

We ran four sessions of the Alternative Advice treatment with 24 subjects in each session. All sessions were intended to run for a total of 32 rounds, although one was terminated after 19 rounds due to a software failure. The 32 rounds were subdivided in two cycles of twelve rounds and one cycle of eight rounds, matching the two cycles of an inexperienced subject session and the first cycle of an experienced subject session in all of the other experiments reported on. Given the high levels of strategic play observed in the latter part of experienced subject sessions, we felt that little would be learned by running longer sessions or bringing subjects back for an experienced subject sessions. Eschewing experienced subject sessions also eliminated any difficulties with keeping subjects in the same role, a key consideration. As in other sessions, subjects were randomly assigned teammates and roles, and were flipped between the M and E roles after each half cycle.

[Insert Fig 6 here]

The main result for the Alternative Advice treatment can be seen in Figure 6. This shows the percentage of strategic play by MLs, broken down by cycle. As points of comparison, strategic play is also shown for the 1x1 treatment and advises from the Advice treatment. While advisees from the Advice treatment consistently do better than the 1x1 treatment, advisees from the Alternative Advice treatment consistently do worse than the 1 x 1 treatment. However, this difference is not especially large.

To determine whether these differences are significant, we ran regressions similar to those in Table 4. The dataset combines data from the Alternative Advice treatment, data from the 1 x 1 treatment (pooled with self-advice data), and data from advisees in the Advice treatment. As with the regression reported in Table 4, this is a logit model with random effects nested at the session and individual levels and a control for the entry rate differential. The base is the Alternative Advice treatment, so treatment effects are identified relative to this treatment for the same cycle. Controls are included, by cycle, for the 1x1 and Alternative Advice treatments. Table 8 reports the results of this exercise. The first column reports parameters for the 1x1 treatment controls and the second column gives parameters for the advisee controls. The estimated difference between the Alternative Advice and 1 x 1 treatments is positive (i.e. there is *more* strategic play in the 1 x 1 treatment), but only significant in the first cycle of inexperienced play. That is, the Alternative Advice treatment does not improve advisees level of strategic play compared the 1-1 treatment where agents choose strictly on their own. This contrasts with the original advice treatment where advisees outperform the 1-1 treatment, significantly so at times. The differences between advisees from the Advice treatment and the Alternative Advice treatment are significant in all three cycles. The original advice treatment, where advisors are playing the game and only receive a fraction of the advisees' payoffs, leads to unambiguously more strategic play by advisees than the new advice treatment where advisors are not playing and receive the same payoff as advisees.

[Insert Table 8 here]

Underlying this difference in performance between the Advice and Alternative Advice treatments are differences in the advice sent. There was no shortage of advice in the Alternative Advice treatment. Every advisee received relevant advice (i.e. advice which related to how the game should be played) at least once, and they averaged 16.5 pieces of relevant advice. Even accounting for the different number of periods, this is quite a bit more relevant advice than in the Advice treatment. The problem is that the quality of this advice was poor. In inexperienced subject sessions of the Advice treatment, 46% of advisees were told to play strategically as an ML at least once. This figure drops to 27% in the Alternative Advice treatment even though they had 32 periods to receive advice as opposed to only 24 periods in an inexperienced subject session of the Advice treatment.

Drilling deeper, we can get a sense of why advisors offered less useful advice in the Alternative Advice treatment. In this treatment, advisees generally either received advice to play strategically starting in the first cycle of inexperienced play (21%) or never received it (73%). This differs from the original Advice treatment where inexperienced advisees were just as likely to initially receive advice to play strategically in Cycle 2 (23%) as Cycle 1 (23%). On a more qualitative level, while the frequency of offering advice changes little over time in the Alternative Advice treatment, the messages become terse repetitions of the same advice, suggesting little effort on the part of advisees. In the original Advice treatment it is clear that the advisors learned to play strategically over time as they actually played the game, becoming increasingly likely to give good advice as a consequence. In the Alternative Advice treatment advisors do not actually play the game and there is no increase in the number of advisors providing sound advice on strategic play. One possible explanation for this difference is that advisors became disengaged, facing the same situation repeatedly without the ability to take any direct actions. That is, even though advisors in the original Advice treatment had less incentive to give advice, and more things to do other than giving advice, having them play the game served the critical role of keeping them engaged in the task, continuing to learn to play strategically, and passing this on as a consequence.

It's been suggested that the results of our Alternative Advice contradict the work of Merlo and Schotter (1999, 2003) in which observers (who strictly observed in silence) performed better than the agents they were observing. But in this case observers were told at the start of the session that what they were observing was relevant to a high stakes decision they would eventually make, which presumably focused their attention on the problem at hand.[39] This is consistent with the idea that having to make a decision, or anticipating having to make a decision, is a critical element in the good performance of observers.

Finally, it is not clear that the poor quality of advice in the Alternative Advice treatment speaks to the value of advisors used by firms outside the lab. Even though consultants do not typically make decisions for themselves, any tendency toward social loafing is likely disciplined by strong reputational reasons to remain engaged.[40] We rather think the main takeaway is methodological. In experimental economics we are taught the value of making treatments as parallel as possible, but sometimes this is not the case. While there is obvious harm in having an experimental design with a confound, in this case subject boredom may have overcome the benefits of parallelism.

[39] In these experiments subjects made a large number (75) of individual choices in an optimization task. Merlo and Schotter conjecture that the reason the original decision makers performed worse than the observers is that the small repeated payoffs in the initial treatment focused attention on the stimulus-aspect of the problem, preventing subjects from learning the tradeoffs involved in the question at hand. The observers, not receiving any payoffs, did not get caught up in this. Our task is quite different, focusing on strategic reasoning rather than the ability to solve an optimization problem. See Iyengar and Schotter (2008) for results related to those of Merlo and Schotter.

[40] Consultants are often hired for their prior expertise with the question at hand and learning is not an issue per se.

Conclusion 8: The Alternative Advice treatment does nothing to change our conclusion that bilateral communication is a necessary ingredient for the strong performance of teams in the limit pricing game.

VII. Summary and Discussion: The primary motivation for the present paper was to better understand how two person teams are able to meet and beat the truth wins (TW) benchmark for strategic play in a signaling game. With this goal in mind we introduced two new treatments into the limit pricing game: (1) a self-advice treatment in which players were encouraged to express to write down their thoughts "regarding what was going on" and (2) an advisor-advisee treatment with one-way communications between advisors and their advisee. The self-advice treatment had essentially no impact compared to individual subject play without self-advice. While team play forces individuals to articulate their thought processes, this alone is not sufficient to generate a significant increase in strategic play. The advisor-advisee relationship resulted in increased levels of strategic play for advisees compared to individuals, but fell well short of their (modified) TW benchmark. This significant increase in the level of strategic play for advisees is consistent with the idea that two heads are better than one. But the failure to meet the modified TW benchmark indicates a significant loss of information in the advisor-advisee relationship relative to two person teams acting jointly with continuous chat. This implies that bilateral communication is a necessary condition for meeting and beating the TW benchmark, even though teams are susceptible to free riding and/or idle chatter that distracts from dealing with the problem at hand.

The failure of advisees to meet the TW benchmark results from (i) advisors who act strategically but fail to pass the relevant advice and (ii) advisees who receive sound advice but fail to heed it. Although the latter has been well documented in the psychology and economics literature on advice giving, the former has not been reported previously. Looking at the team dialogues suggests one factor behind the failure to follow sound advice – when advisees do not understand the advice, one way communication prevents advisors from identifying and correcting any misunderstandings.

The failure of advisors who play strategically to advise their partners to do so cannot be attributed to free riding given the incentive structure of the experiments, and cannot be explained by hedging or the level of experience advisors had with strategic play. Rather, gender differences were, to our surprise, playing a significant role. Male advisors who played strategically consistently passed on relevant advice more often than female advisors. We believe this result is related to well established differences in confidence levels between men and women in a number of strategic situations (Niederle and Vesterlund, 2007). Given the inevitable variation in confidence levels *within* the genders, this also implies that between subject differences in confidence levels account for much of the within gender variation in passing on advice. An interesting avenue for future research is experiments that combine the advice treatment with direct elicitation of advisors' confidence in their ability to play the game. We conjecture that the gender effect will largely vanish once we control for differing levels of confidence.

One particularly surprising feature of our data is that the positive impact of advice on strategic play is essentially independent of whether or not a sound explanation accompanies the advice to take a new, and improved, course of action. This suggests that the main impact of advice is to simply get advisees to consider some new options. This contradicts economists' usual assumption that agents carefully consider all the options at their disposal prior to making their choices, and (in the spirit of models of rational inattention) indicates that learning models should not only account for learning over a fixed set of available options but should also allow for exploration of the set of available options where some are considered and abandoned while others are only evaluated when some external prompt is provided. Experiments using techniques such as mouse-lab or eye-tracking would be helpful in directly confirming this.

The relative success of freely interacting teams compared to advisees bears some important similarities to the difference between successful and unsuccessful consulting relationships in field settings: Client organizations perceive effective consulting as a two-way interaction with knowledge outflow from client to consultant influencing the quantity and quality of knowledge inflow from the consultant to the client organization (Todorova, 2004, Newell, 2005). In this respect teams, with their two way interaction, strongly outperform "clients" in the advice treatment even though there are two heads to work with in both cases. This may also have implications for successful discussions between different divisions within a hierarchal organization, particularly to the extent that lower ranked members might be constrained in their discussions and/or interactions with higher ranked members.

Finally, it is worth discussing why teams are able to meet and beat the TW benchmark in this signaling game, but fail in other environments where there is a demonstrably correct answer.[41] In this respect, the contrast with the takeover game (Bazerman and Samuelson, 1983) where teams typically fail to meet the TW benchmark (Casari, Zhang, and Jackson, forthcoming; Cooper and Sutter, 2015), is instructive. In that environment, it is one thing to become aware of (either consciously or subconsciously) and respond to the adverse selection effect. However, it is an entirely different thing to be able to calculate the precise response to this insight, which requires some training in statistics among other things. In this case teams have little advantage over individuals, since teammates cannot communicate relevant insights that they don't have. In contrast, in our signaling game it's clear relatively early on that the low cost monopolists want to distinguish themselves from the high cost types, and that choice of 5, 6 or 7 (particularly the 6 and 7 with their negative payoff for high cost types) will do the trick; the insight is immediate and requires minimal computations regarding the tradeoff between sticking with 4 and choosing to limit price. In this respect it will be interesting in future work to see how well teams perform compared to individuals in more subtle signaling games.

References

Bazerman, Max and Samuelson, William. "I won the auction but don't want the prize." *Journal of Conflict Resolution*, 1983, 27, pp 618-634.

Bonaccio, Silvia and Dalal, Reeshad S. "Advice Taking and Decision-Making: An Integrative Review of the Literature." *Organizational Behavior and Human Decision Processes*, 2006, 101, pp. 127-151.

Casari, Marco, Zhang, Jingling, and Jackson, Christine. "When do groups perfume better than individuals? An acquiring a company experiment," *Experimental Economics*, forthcoming.

Charness, Gary and Sutter, Matthias. "Groups Male Better Self-Interested Decisions." *Journal of Economic Perspectives*, 2012, 26 (3), pp. 157-176.

Chaudhuri, Ananish, Schotter, Andrew, and Sopher, Barry. "Learning in tournaments with inter-generational advice," *Economics Bulletin*, 2006, 3 (26), pp. 1-16.

Cooper, David J. and Kagel, John H. "Are Two Heads Better Than One? Team versus Individual Play in Signaling Games." *American Economic Review*, 2005, 95(3), pp. 477-509.

[41] For discussion of demonstrability requirements in the psychology literature see Laughlin and Ellis (1986).

_________ and ___________. ““The Role of Context and Team Play in Cross-Game Learning” *Journal of the European Economic Association*, 2009, 7(5), pp. 1101-1139.

__________ and _____________. “The Effects of Cognitive Ability on the Development of Strategic Play in Signaling Games.” (work in progress).

__________ and **Sutter, Matthias**. “Endogenous role assignment and team performance.” 2015, working paper.

Davis, James, H. “Some Compelling Intuitions About Group Consensus Decisions, Theoretical and Empirical Research, and Interpersonal Aggregation Phenomena: Selected Examples, 1950-1990.” *Organizational Behavior and Human Decision Processes,* June 1992, *52*(1), pp. 3-38.

Kugler, Tamara, Kausel, E. and Kocher, Martin. “Are Groups More Rational than Individuals? A review of Interactive Decision Making in Groups,” *Wiley Interdisciplinary Reviews: Cognitive Science*, 2012, 3 (4), pp. 471-482.

Iyengar, Raghuram and **Schotter, Andrew** “Learning with a Meddlesome Boss: An Experiment in Sloppy Principal-Agent Relationships.” *Experimental Economics,* 2008, 11, pp 154-73.

Laughlin, Patrick R. and Ellis, Alan L. “Demonstrability and social combination processes on mathematical intellective tasks,” *Journal of Experimental Social Psychology,* 1986, 22, pp. 177-189.

Lerner, Jennifer S. and Tetlock, Philip E. “Accounting for the Effects of Accountability.” *Psychological Bulletin,* 1999, 125(2), pp. 255-275.

Lord, Charles G., Lepper, Mark R., and Preston, Elizabeth. “Considering the opposite: a corrective strategy for social judgment,” *Journal of Personality and Social Psychology,*1984, 47 (6), ppp. 1231-1243.

Lorge, Irving and Solomon, Herbert. “Two models of group behavior in the solution of eureka-type problems.” *Psychometrika*, June 1955, *20*(2), pp. 139-48.

Milgrom, Paul. and Roberts, John. “Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis.” *Econometrica*, 1982, 50(2), pp. 443-59.

Mussweiler, Thomas, Strack, Fritz, and Pfeiffer, Tim. “Overcoming the inevitable anchoring effect: Considering the opposite compensates for selective accessibility,” *Personality and Social Psychology Bulletin*, 26 (9), pp. 1142-1150.

Newell, Susan. “The Fallacy of Simplistic Notions of the Transfer of ‘Best Practice.” In: Anthony F. Buono and Flemming Poulfelt (Eds.), *Challenges and Issues in Knowledge Management* (Research in Management Consulting,) 2005, 5. Greenwich, CT: Information Age Publishing.

Niederle, Muriel and Vesterlund, Lise “Do Women Shy Away from Competition? Do Men Compete Too Much” *Quarterly Journal of Economics,* 2007, 122 (3), pp. 1067-1101.

Schotter, Andrew. “Decision Making with Naïve Advice.” *American Economic Review***,** 2003, 93(2), pp. 196-201.

_________ and Sopher, Barry. “Social Learning and Coordination Convnetionsin Intergenerational Games: An Experimental Study.” *Journal of Political Economy,* 2003, 111 (3), pp. 498- 524.

Todorova, Gergana. “Exploring Knowledge Issues in the Consulting Relationship.” 2004, Information Age Publishing, USA, in Creative Counseling: Innovative Perspectives on Management Consulting, (ed) Anthony Bruno. pp. 73-98.

Figure 1: Play in the 1 x 1 Treatment

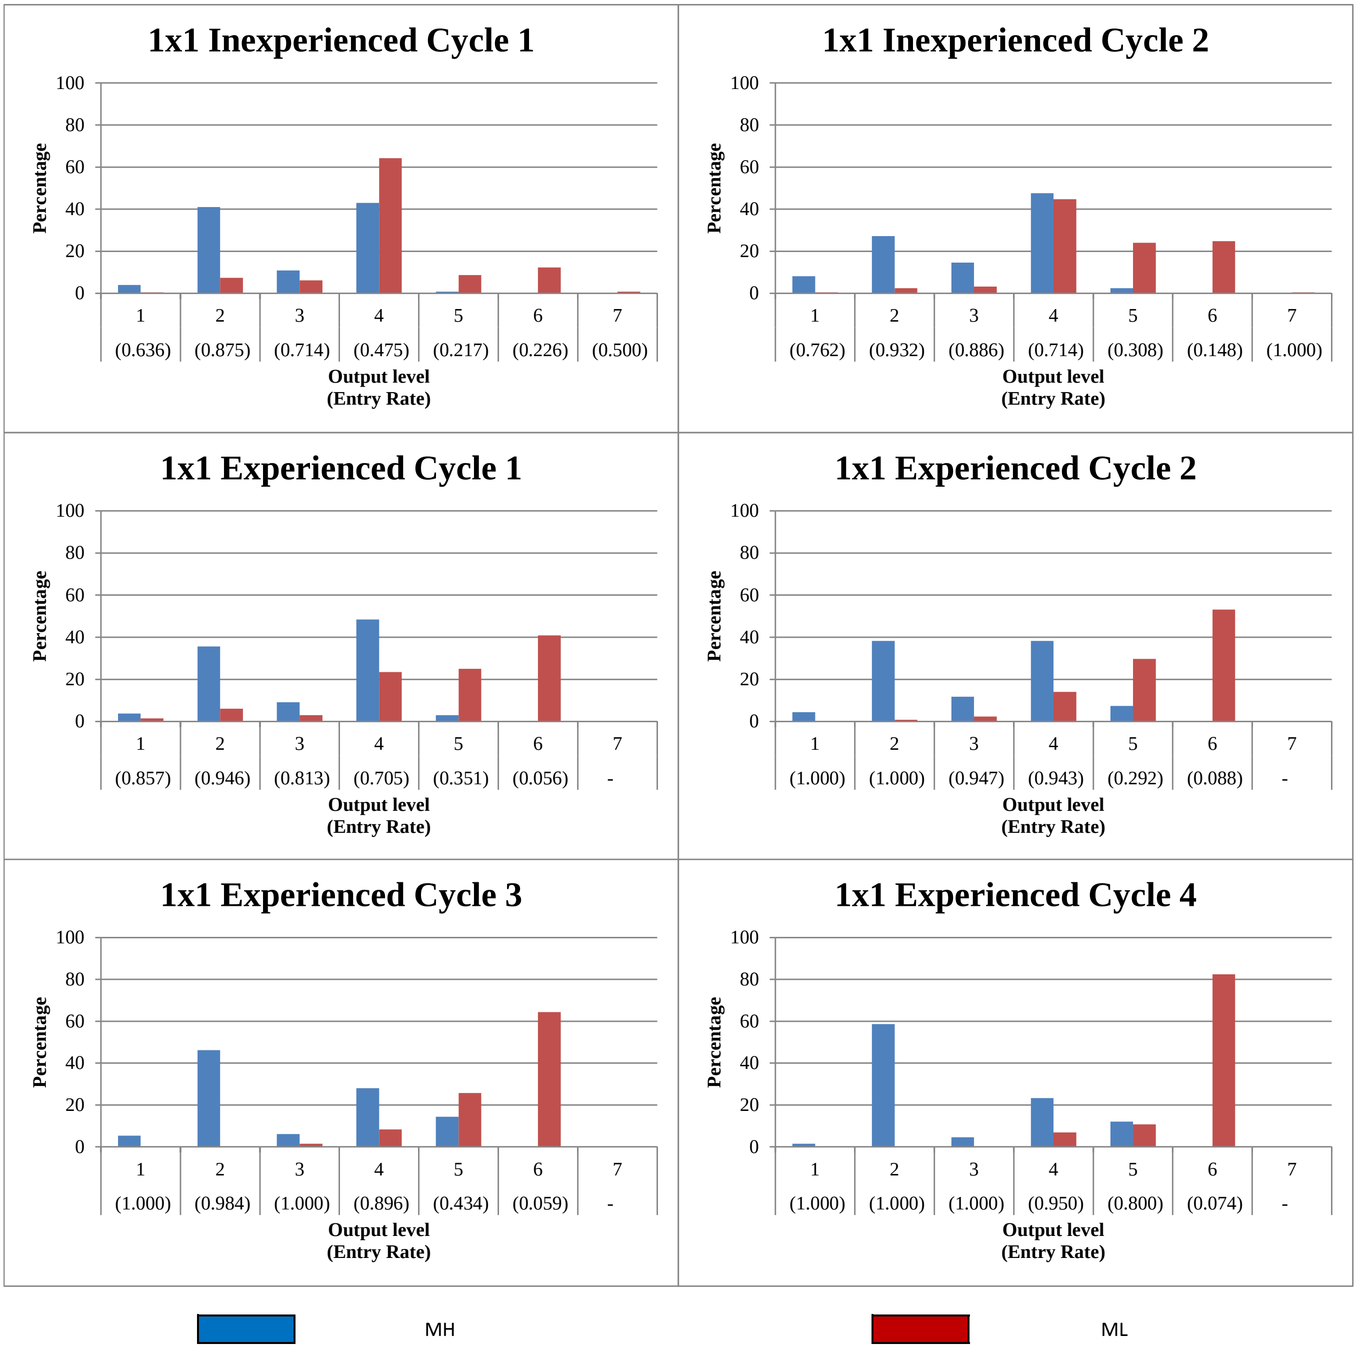

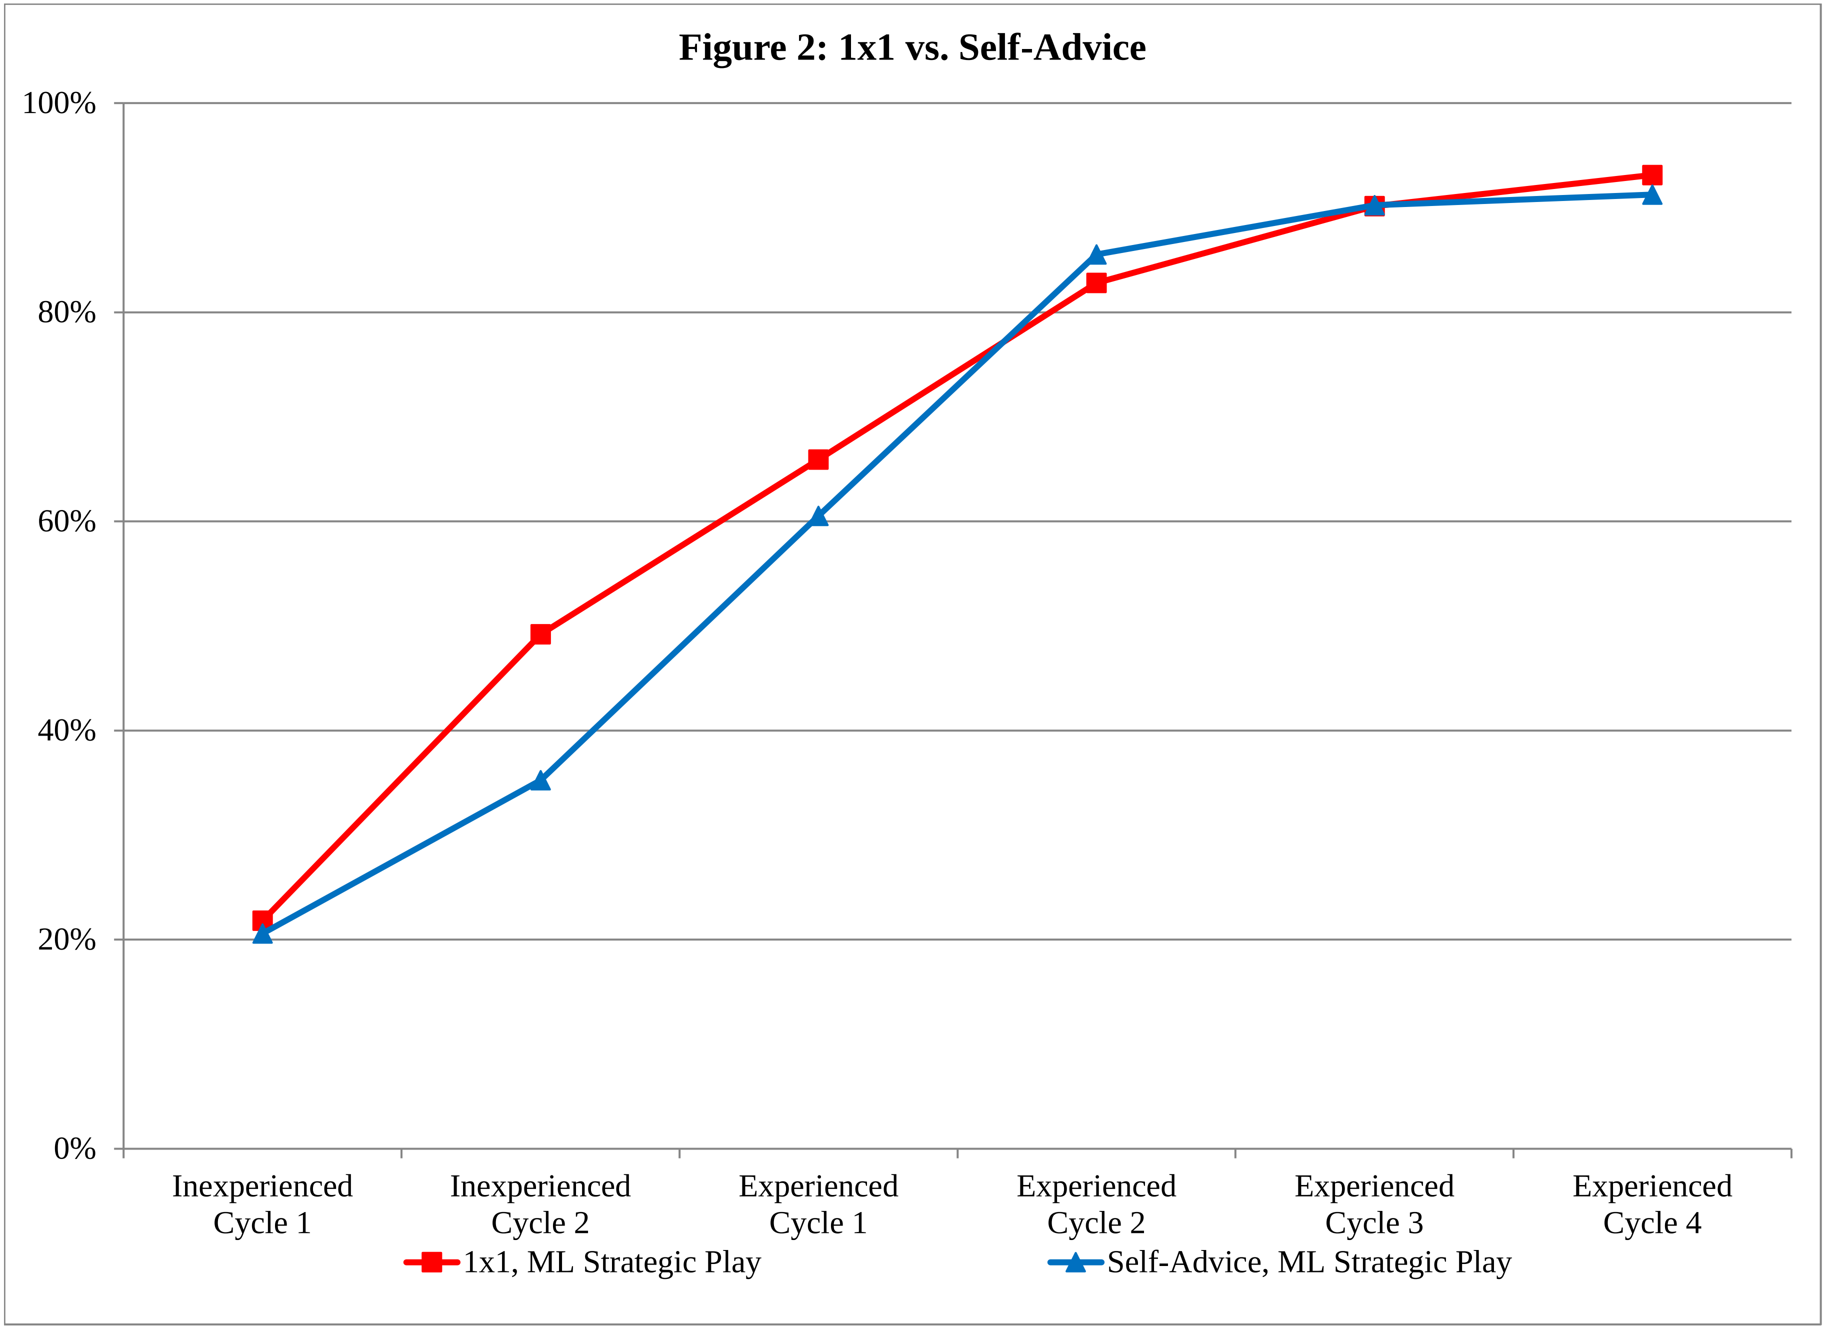
Figure 2: 1x1 vs. Self-Advice
100%
80%
60%
40%
20%
0%
Inexperienced Cycle 1
Inexperienced Cycle 2
Experienced Cycle 1
Experienced Cycle 2
Experienced Cycle 3
Experienced Cycle 4
1x1, ML Strategic Play
Self-Advice, ML Strategic Play

Figure 4: Strategic Play by MLs, 1x1 vs. 2x2

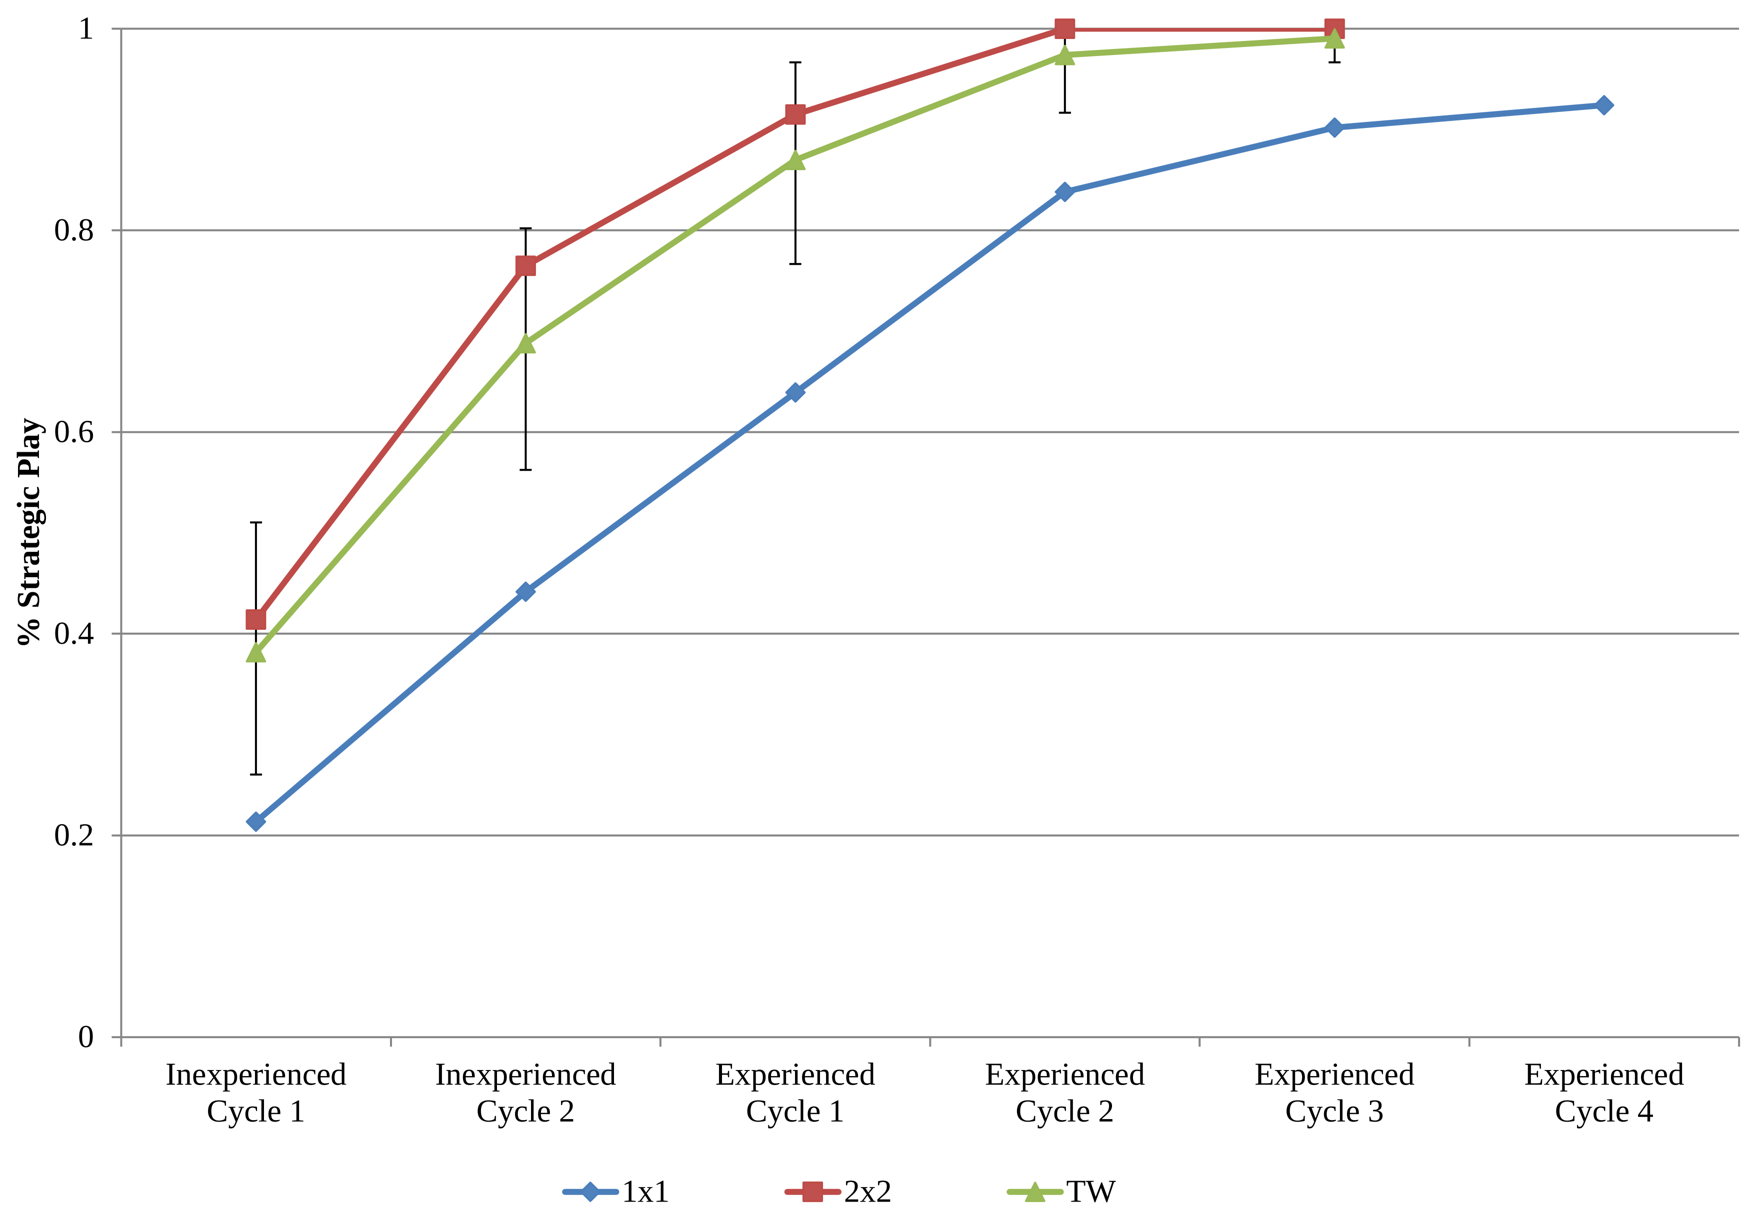

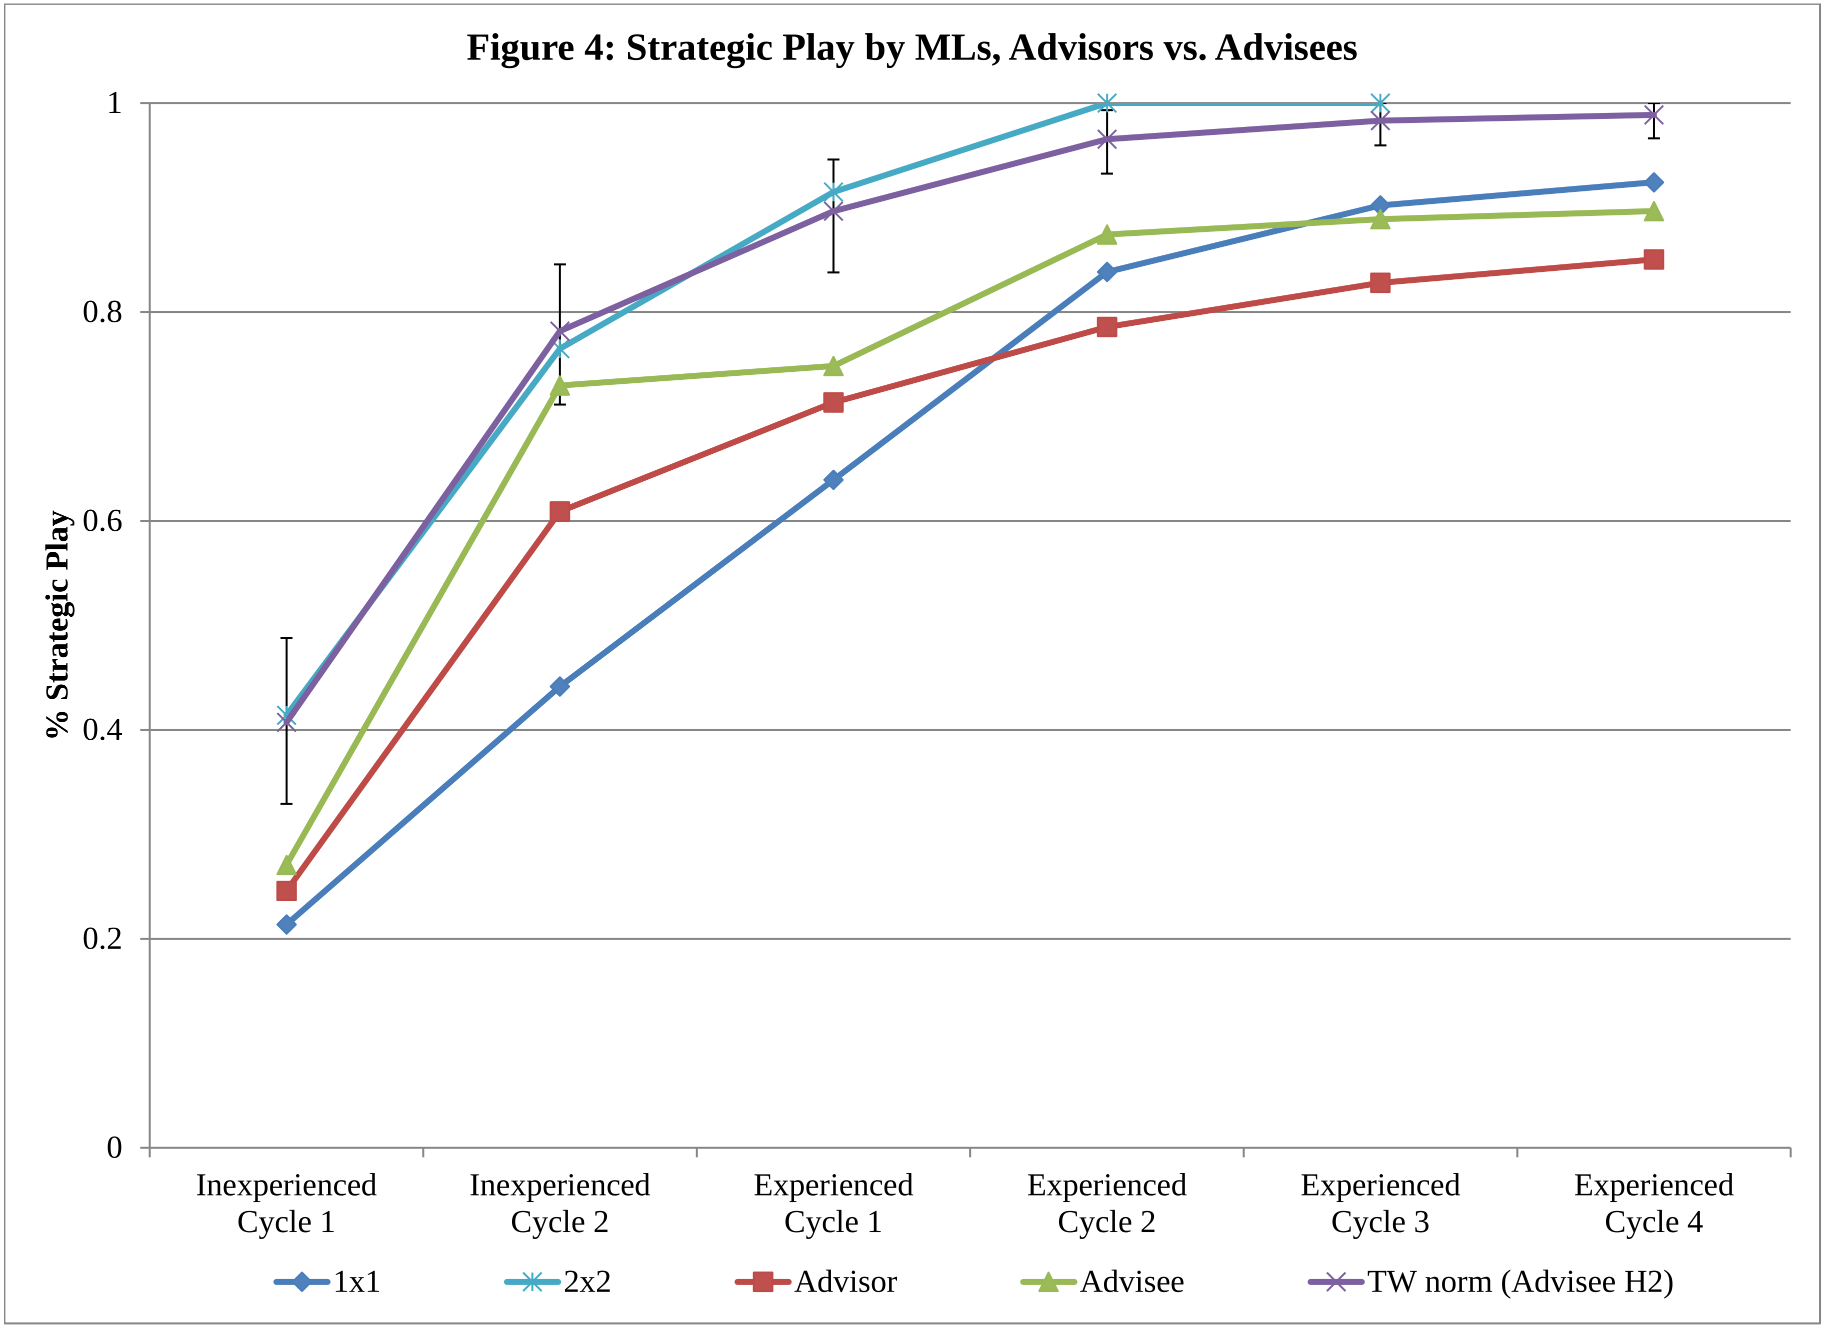
Figure 4: Strategic Play by MLs, Advisors vs. Advisees
% Strategic Play
1
0.8
0.6
0.4
0.2
0
Inexperienced Cycle 1
Inexperienced Cycle 2
Experienced Cycle 1
Experienced Cycle 2
Experienced Cycle 3
Experienced Cycle 4
1x1
2x2
Advisor
Advisee
TW norm (Advisee H2)

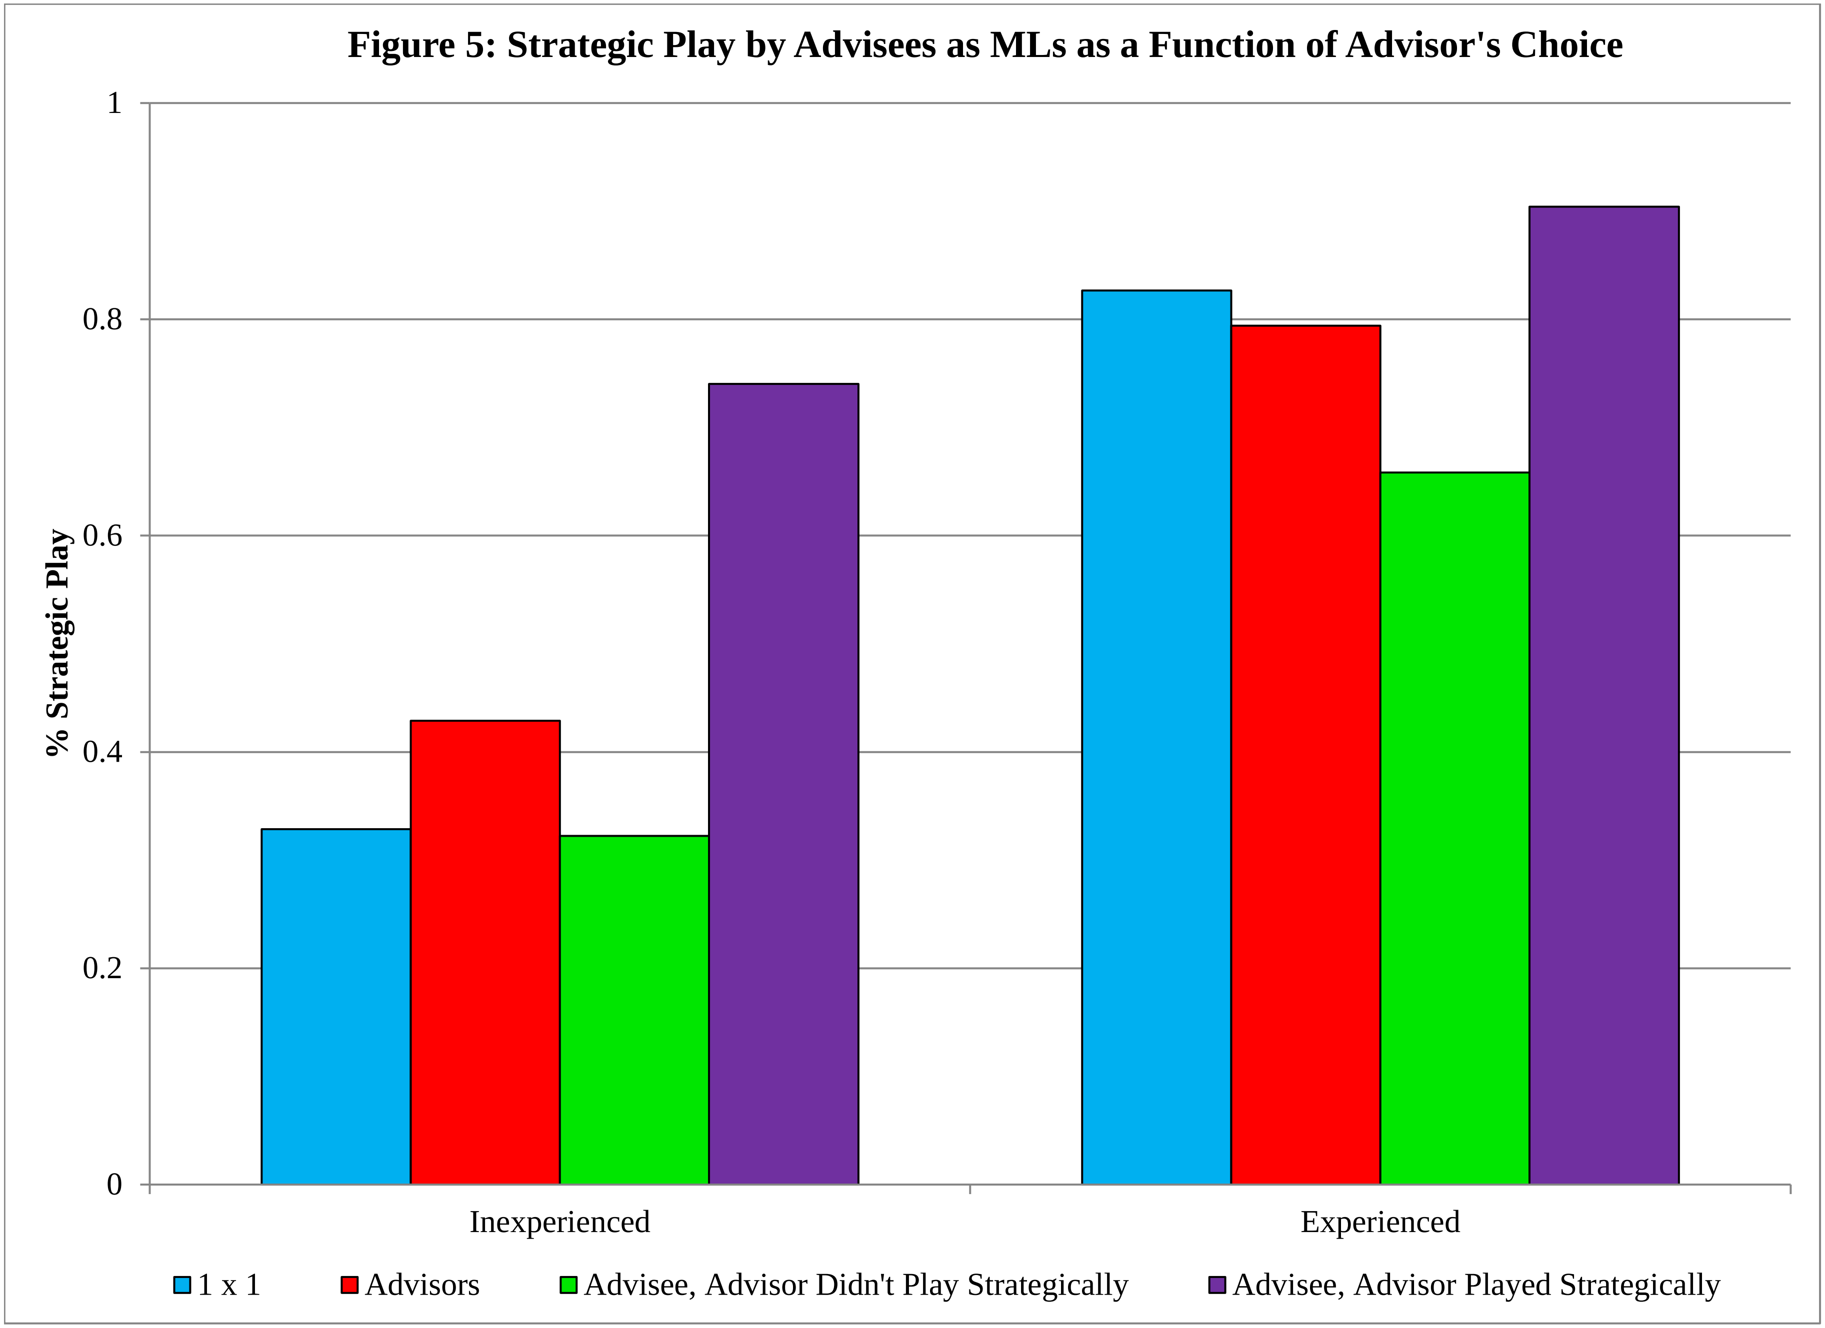
Figure 5: Strategic Play by Advisees as MLs as a Function of Advisor's Choice
1
0.8
0.6
0.4
0.2
0
% Strategic Play
Inexperienced
Experienced
1 x 1
Advisors
Advisee, Advisor Didn't Play Strategically
Advisee, Advisor Played Strategically

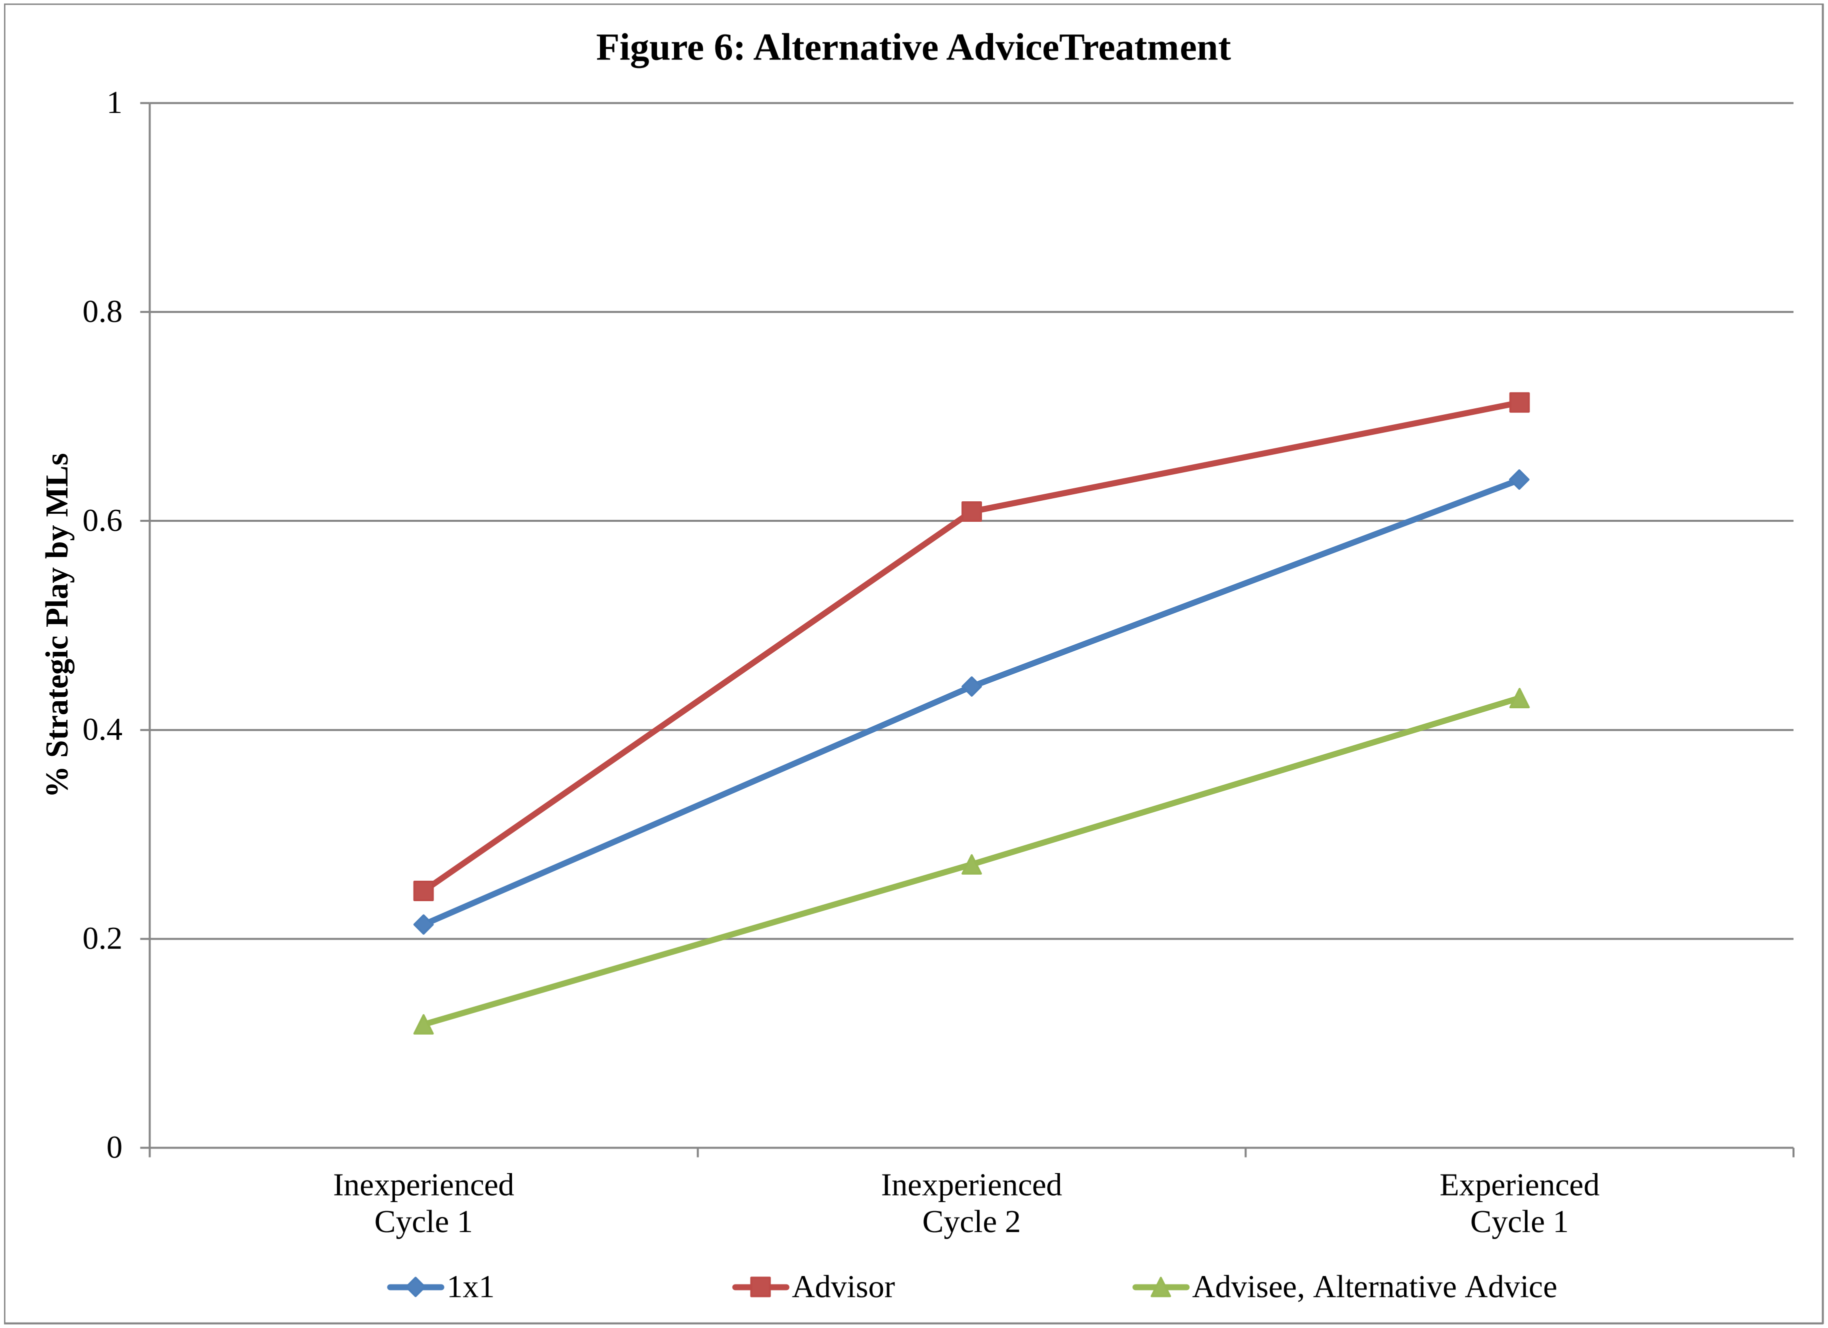
Figure 6: Alternative AdviceTreatment
1
0.8
0.6
0.4
0.2
0
% Strategic Play by MLs
Inexperienced
Cycle 1
Inexperienced
Cycle 2
Experienced
Cycle 1
1x1
Advisor
Advisee, Alternative Advice

Table 1a: Monopolist Payoffs

High Cost Monopolist (MH)		
Monopolist Action	Entrant's Response	
	IN	OUT
1	150	426
2	168	444
3	150	426
4	132	408
5	56	182
6	-188	-38
7	-292	-126

Low Cost Monopolist (ML)		
Monopolist Action	Entrant's Response	
	IN	OUT
1	250	542
2	276	568
3	330	606
4	352	628
5	334	610
6	316	592
7	213	486

Table 1b: Entrant Payoffs

Entrant's Strategy	Monopolist's Type		Expected Payoff
	High Cost	Low Cost	
IN	500	200	350
OUT	250	250	250

Notes: The highlighting of cells and payoffs is for expositional purposes – subjects' payoff tables did not include any highlighted material. Information on expected payoffs (highlighted in blue) was not displayed on subjects' payoff tables.

Table 2: Summary of Sessions

	Inexperienced	Experienced
1 x 1	6 Sessions 82 Subjects	5 Sessions 66 Subjects
2 x 2 (Team)	3 Sessions 64 Subjects	3 Sessions 59 Subjects
Advisor/Advisee	9 Sessions 163 Subjects	9 Sessions 146 Subjects
Self-Advice	3 Sessions 46 Subjects	3 Sessions 40 Subjects

Notes: In advisor/advisee sessions that were one subject short, an undergraduate RA played as the advisee so no subjects would need to be discarded.

Table 3: Effect of Self Advice

Variable	Parameter Estimate	Standard Error
Inexperienced Cycle 2	2.380***	.398
Experienced Cycle 1	4.173***	.868
Experienced Cycle 2	6.258***	1.007
Experienced Cycle 3	7.099***	1.005
Experienced Cycle 4	8.853***	1.177
Self-advice * Inexperienced Cycle 1	.040	.944
Self-advice * Inexperienced Cycle 2	-1.110	.891
Self-advice * Experienced Cycle 1	-.812	.971
Self-advice * Experienced Cycle 2	.315	1.096
Self-advice * Experienced Cycle 3	-.482	1.149
Self-advice * Experienced Cycle 4	-1.746	1.238
Number of Subjects	-.167	.203
Entry Rate Differential	1.500*	.808
Log-likelihood	-551.08	

Notes: The regression is based on 1601 observations from 129 individuals in the 17 sessions of the 1x1 and self-advice treatments, and includes random effects at the session and individual levels. One (*), two (**), and three (***) stars indicate statistical significance at the 10%, 5%, and 1% levels respectively.

Table 4: Logit with Nested Random Effects: Treatment Effects on Strategic Play Over Time

Variable	Parameter Estimate	Standard Error
Inexperienced Cycle 2	1.582^{***}	.313
Experienced Cycle 1	3.094^{***}	.567
Experienced Cycle 2	5.243^{***}	.671
Experienced Cycle 3	5.649^{***}	.713
Experienced Cycle 4	6.706^{***}	.766
2x2 * Inexperienced Cycle 1	1.721^{**}	.757
2x2 * Inexperienced Cycle 2	3.077^{***}	.816
2x2 * Experienced Cycle 1	2.603^{**}	1.060
2x2 * Experienced Cycle 2	Dropped, 100% Strategic Play	
2x2 * Experienced Cycle 3	Dropped, 100% Strategic Play	
Advisor * Inexperienced Cycle 1	.423	.602
Advisor * Inexperienced Cycle 2	1.155^{**}	.571
Advisor * Experienced Cycle 1	.638	.623
Advisor * Experienced Cycle 2	-.779	.672
Advisor * Experienced Cycle 3	-.870	.706
Advisor * Experienced Cycle 4	-1.571^{**}	.748
Advisee * Inexperienced Cycle 1	.630	.600
Advisee * Inexperienced Cycle 2	2.158^{***}	.579
Advisee * Experienced Cycle 1	1.053^{*}	.631
Advisee * Experienced Cycle 2	.385	.697
Advisee * Experienced Cycle 3	.154	.736
Advisee * Experienced Cycle 4	-.812	.773
Number of Players	-.003	.046
Entry Rate Differential	2.105^{***}	.538
Log-likelihood	-1472.52	

Notes: The regression includes 3993 observations from 337 individuals/teams in 41 sessions, and includes random effects at the session and individual levels. One (*), two (**), and three (***) stars indicate statistical significance at the 10%, 5%, and 1% levels respectively.

Table 5: Hazard Models, Giving Advice to Play Strategically

	Model 1	Model 2	Model 3
Gender (0 = male, 1 = female)	-.535** (.251)	-.477* (.270)	1.248*** (.484)
Experienced	.150 (.400)	.109 (.445)	-.084 (.464)
Number of Advisors In Session	.028 (0.54)	.057 (.058)	.054 (.059)
Entry Rate Differential	.721 (.692)	-.571 (.801)	-.352 (.840)
% Strategic Play		1.987*** (.380)	3.092*** (.491)
% Strategic Play * Gender			-2.350*** (.661)
Log-likelihood	-368.77	-354.25	-349.16

Notes: All regressions include 158 observations from 120 advisors in 18 sessions. Standard errors (reported in parentheses) are corrected for clustering at the advisor level. One (*), two (**), and three (***) stars indicate statistical significance at the 10%, 5%, and 1% levels respectively.

Table 6: Strategic Play as an ML as a Function of Type of Advice Received

Inexperienced Subjects

		Given Explanation for Strategic Play as an ML	
		No	Yes
Advised to Play Strategically as an ML	No	.423 (319)	.500 (10)
	Yes	.655 (58)	.670 (97)

Experienced Subjects

		Given Explanation for Strategic Play as an ML	
		No	Yes
Advised to Play Strategically as an ML	No	.783 (258)	N/A (0)
	Yes	.812 (101)	.951 (224)

Table 7: Logit with Nested Random Effects: Effect of Receiving Advice

Variable	Parameter Estimate	Standard Error
Inexperienced Cycle 2	3.173***	.571
Experienced Cycle 1	2.366**	1.237
Experienced Cycle 2	3.973***	1.324
Experienced Cycle 3	3.966***	1.357
Experienced Cycle 4	4.459***	1.405
Inexperienced, Received Advice to Play Strategically	2.480**	.972
Inexperienced, Received Explanation for Strategic Play	.019	1.201
Inexperienced, Received Bad Advice	-1.416**	.660
Experienced, Received Advice to Play Strategically	2.621*	1.403
Experienced, Received Explanation for Strategic Play	.294	1.536
Experienced, Received Bad Advice	.992	.842
Number of Players	-.124	.200
Entry Rate Differential	5.012***	1.549
Log-likelihood	-325.54	

Notes: All regressions include fixed effects for the current cycle as well as random effects at the session and individual levels. Models 1 and 2 include 1067 observations from 18 sessions and 124 advisees. Model 3 includes 1039 observations from 18 sessions and 121 advisees. One (*), two (**), and three (***) stars indicate statistical significance at the 10%, 5%, and 1% levels respectively.

Table 8: Probit, Treatment Effect of Alternative Advice Treatment

Cycle	Alt. Advice vs. 1x1	Alt. Advice vs. Advisees
Inexperienced Cycle 1	-1.878* (0.973)	-2.198* (1.158)
Inexperienced Cycle 2	-1.048 (0.908)	-3.777*** (1.109)
Experienced Cycle 1	-1.562 (0.974)	-2.814** (1.104)

Notes: One (*), two (**), and three (***) stars indicate statistical significance at the 10%, 5%, and 1% levels respectively.