

David J. Cooper · John H. Kagel

Learning and transfer in signaling games

Received: 22 November 2004 / Revised: 17 November 2006 /

Published online: 16 January 2007

© Springer-Verlag 2007

Abstract We explore how learning to play strategically in one signaling game promotes strategic play in a related signaling game. Following convergence to a pooling equilibrium, payoffs are changed to only support separating equilibria. More strategic play is observed following the change in payoffs than for inexperienced subjects in control sessions, contrary to the prediction of a fictitious play learning model. Introducing a growing proportion of sophisticated learners, subjects who anticipate responders' behavior following the change in payoffs, enables the model to capture the positive cross-game learning observed in the data.

Keywords Learning · Learning transfer · Cross-game learning · Signaling games · Experiment

JEL Classification Numbers C72 · C92 · D82 · L12

Research support from the National Science Foundation grant number SBR9809538 is gratefully acknowledged. We have received research support from Jo Ducey, Guillaume Frechette, Steve Lehrer, and Carol Kraker Stockman. We have benefitted from comments of Eric Bettinger, John Ham, Jim Rebeitzer, Bob Slonim and seminar participants at Case Western Reserve University, Ohio State University, the University of Mississippi, the University of Illinois, and Purdue University. The usual caveat applies.

D. J. Cooper (✉)

Department of Economics, Weatherhead School of Management,
Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106-7206, USA
E-mail: david.cooper@case.edu
Tel.: +1-216-3684294
Fax.: +1-216-3685039

J. H. Kagel

Department of Economics, 410 Arps Hall, Ohio State University, 1945 North High Street,
Columbus, OH 43210-1172, USA
E-mail: kagel@economics.sbs.ohio-state.edu

1 Introduction

Numerous experiments show that models of learning in which players have bounded rationality and only gradually learn how to best respond in a game capture important features in experimental data missed by standard equilibrium approaches based on full rationality of players.¹ Over time subjects generally learn to make decisions that are more consistent with rationality, implying that the stability properties of learning models can serve as a viable foundation for equilibrium analysis. However, virtually all papers on learning employ an environment in which learning takes place within a stationary environment while in many real world settings the game being played changes over time. Stability properties derived in stable environments may be irrelevant if changes in the game disrupt the learning process. The ability to take what has been learned in one game and apply it in another related game is therefore an integral but largely unexplored aspect of learning in games.

The extensive psychology literature on transfer indicates that the ability to generalize across games cannot be taken for granted. Positive transfer usually fails except in settings that are perceived as being quite similar. This failure follows in part from subjects' inability to recognize underlying concepts that allow them to generalize between settings (Gick and Holyoak 1980; Perkins and Salomon 1988; Salomon and Perkins 1989). While suggestive, the direct relevance of these findings for economic games is questionable. Psychology studies of learning transfer tend to be one-shot in nature, both in terms of what was initially learned and in terms of the new learning environment. In contrast cross game learning issues in economics are largely concerned with whether having adjusted *over time* to equilibrium in one game will speed up the adjustment *over time* to a new equilibrium in a related game. Additionally, the insights gained in many psychology studies are algorithmic in nature (e.g., what is the best method of solving a logic problem), while successful play in many games revolves around psychological insights (e.g., is my opponent trying to fool me). This mismatch between studies of learning transfer in psychology and game theoretic settings underlines the need to study cross game learning.

The goals of our experiment are to determine whether or not positive transfer occurs between related games and, more importantly, to identify the mechanism by which transfer occurs.² We study these issues in the context of a well-known signaling game from the industrial organization literature, Milgrom and Roberts' (1982) entry limit pricing game. Strategic play in this game revolves around an incumbent monopolist attempting to deter entry by signaling it will be a tough competitor for a potential entrant. Past experiments have found that strategic play only emerges gradually, with most monopolists initially ignoring the strategic implications of their choices on entrants' responses (Cooper et al. 1997a,b).

The limit pricing game provides a rich environment for studying transfer. Like most signaling games, it features multiple equilibria including pure strategy pooling and separating equilibria. This multiplicity allows us to confront subjects with related games that require quite different actions to play strategically. Further,

¹ See Camerer (2003) for a recent review of the existing literature and citations to same.

² For other studies of learning transfer in game theoretic settings see Kagel and Levin (1986) and Ho et al. (1998). The contribution of our work lies not in being the first to consider learning transfer in games but rather in our exploration of the mechanism underlying this transfer.

strategic play is clearly identifiable in the limit pricing game, making it easy to measure the extent to which there is cross-game learning. While our results are derived within the framework of a single game, the insights generated are likely to apply broadly as the main concepts needed to play strategically in the limit pricing game (e.g., think from the other player's point of view, anticipate that others will attempt to glean information from your actions) are also valid in many other games.

We confront subjects with a challenging test of their ability to transfer learning between games. In the initial game, entrants' payoffs support a pure strategy pooling equilibrium to which inexperienced subjects' play reliably converges. In this equilibrium high cost monopolists act strategically, imitating the low cost monopolists. Entrants' payoffs are then changed to eliminate the pooling equilibrium, leaving only pure strategy separating equilibria. While conceptually similar, strategic behavior in the second game requires substantially different actions than in the first game, as it is now the low cost monopolists who must act strategically, distinguishing themselves from high cost monopolists.

Ex ante, neither the psychology literature or the economics literature on learning lead us to expect much cross-game learning following the change in entrants' payoffs. For positive transfer to occur, it is not sufficient that subjects have learned how to play strategically in the initial game. Subjects must also understand why strategic play works in the first game and recognize that similar concepts apply in the second game. It is precisely this ability to use underlying concepts (as opposed to merely continuing use of a previously successful strategy without understanding the reasons for its success) that the psychology literature identifies as a sticking point for positive transfer.

This point can be made formally using a fictitious play learning model that has worked well in tracking play from previous signaling game experiments (Cooper et al. 1997b).³ It predicts that strategic play by low cost monopolists immediately following the change in entrants' payoffs will be *less* frequent than in control sessions (negative transfer), and will remain less than in the controls until behavior converges to an equilibrium outcome. This prediction is based on the unsophisticated learning process underlying fictitious play, a feature it shares with other commonly used learning models such as replicator dynamics (see Fudenberg and Levine 1998, Chap. 3), reinforcement learning (Roth and Erev 1995), and EWA (Camerer and Ho 1999). A fictitious play learner treats his opponents as a fixed statistical distribution rather than forming a model of how his opponents make decisions. Because of this, a fictitious play learner does not anticipate any change in his opponents' play when their payoffs are altered. Lacking any concept of why strategic play works in the initial game, fictitious play learners (as well as other types of unsophisticated learners) are poorly equipped to continue playing strategically when the environment changes.

Contrary to the preceding, in our data low cost monopolists show significantly *more* strategic play immediately following the change in entrants' payoffs than in control sessions (positive transfer). In fact, the play of subjects following the crossover is statistically indistinguishable from experienced subjects in control sessions, suggesting that experience with the pooling equilibrium is an almost perfect

³ The fictitious play model of learning was introduced by Robinson (1951). The version we employ is closely related to the stochastic fictitious play model of Fudenberg and Levine (1995).

substitute for experience in the game where the only pure strategy equilibria are separating.

To capture the rapid jump to strategic play observed in the data, we modify the basic fictitious play model to include the possibility that subjects learn in a sophisticated manner, modeling how their opponents make decisions and thereby anticipating the change in responders' behavior following the change in their payoffs. We also allow for the possibility that subjects can change modes of learning, switching from unsophisticated to sophisticated learning. Fitting this model to the data, we find a statistically significant fraction of sophisticated learners in the population. *Moreover, a significant fraction of subjects switch from unsophisticated to sophisticated learning with experience.* With the addition of a growing number of sophisticated learners, the model tracks the jump in strategic play following the change in entrants' payoffs.

Because the mechanism by which transfer occurs, sophisticated learning, is relevant for many games, our results are like to extend beyond the specific environment being studied. More broadly it is clear that many subjects are not the simple-minded automata envisioned by standard learning models. This is good news for game theory, a central idea of which is that agents will try to anticipate the actions of others and respond accordingly. Our results indicate that good models of learning should allow for the *development* of substantial sophistication on the part of subjects over time.

2 The limit pricing game

The games studied here are based on Milgrom and Roberts' (1982) entry limit pricing model. For our purposes, the industrial organization implications of this model are of secondary importance. We therefore employ a stylized version of the model that focuses on the signaling aspects of the game. This section describes the two versions of the game used in our experiments and derives equilibrium predictions for these games.

2.1 The game

The limit pricing game is played between an incumbent monopolist (M) and a potential entrant (E). The game proceeds as follows: (1) M observes its type, high cost (MH) or low cost (ML). The two types are realized with equal probability with this being common knowledge. (2) M chooses one of seven output levels (quantities). M's payoff, shown in Table 1, is contingent on its type, the output level chosen, and the E's response. (3) E sees M's output, but not M's type, and either enters or stays out. This 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) by Ms. E's payoff depends on M's type and on E's decision, not on M's output choice. As a treatment variable, two different payoff tables, Tables 2 and 3, were used for Es. These represent "high cost" and "low cost" Es, respectively. Only one of these tables was in use at any given time.

Three features of Table 1 capture the main strategic elements confronting Ms: (1) *Ceteris paribus*, Ms are better off if Es choose OUT rather than IN. (2) Reflecting lower marginal costs, MLs generally prefer higher output than MHs. This can

be seen in Ms' payoffs should they ignore the effect of their choices on Es' behavior—MLs would choose output 4 as opposed to 2 for MHs. These choices will be referred to as the Ms' "myopic maxima." (3) Output levels 6 and 7 are dominated strategies for MHs, but not MLs. At these outputs MLs can, in theory, perfectly distinguish themselves from MHs.

For either high or low cost entrants (Tables 2 or 3) it pays to play IN when M is known to be an MH type and to play OUT against an ML type. However, given the 50–50 probability of the different M types, the expected value of OUT is greater than IN for high cost entrants (250 vs. 187) and the expected value of IN is greater than OUT for low cost entrants (350 vs. 250).

2.2 Equilibrium Predictions

For the limit pricing game with high cost Es (Tables 1, 2), there exist multiple pure strategy pooling, as well as separating, equilibria.⁴ Pure strategy pooling equilibria occur at output levels 1–5. For example, consider a pooling equilibrium at output 3. Given the prior probabilities over M's type, E's expected value of OUT is greater than IN so that pooling deters entry. Beliefs that support this equilibrium are that any deviation involves an MH type with sufficiently high probability to induce entry. Given these beliefs, both MHs and MLs achieve higher profits at 3 rather than deviating to their myopic maxima. Similar out of equilibrium beliefs support the other pooling equilibria. *Pooling equilibria at outputs 3–5 involves strategic play by MHs as they choose higher output levels than would be optimal if they ignored the impact of their choice on E's response.*

Two pure strategy separating equilibria also exist. In both of these MHs choose output level 2 and are always entered on; MLs either always choose output level

Table 1 Monopolist payoffs

High cost monopolist (MH)			Low cost monopolist (ML)		
Monopolist action	Entrant response		Monopolist action	Entrant response	
	IN	OUT		IN	OUT
1	150	426	1	250	542
2	168	444	2	276	568
3	150	426	3	330	606
4	132	408	4	352	628
5	56	182	5	334	610
6	-188	-38	6	316	592
7	-292	-126	7	213	486

Table 2 Entrant payoffs, high cost entrants

Entrant's strategy	Monopolist's type	
	High cost	Low cost
IN	300	74
OUT	250	250

⁴ All of the equilibria to be described are sequential (Kreps and Wilson 1982).

Table 3 Entrant payoffs, low cost entrants

Entrant's strategy	Monopolist's type	
	High cost	Low cost
IN	500	200
OUT	250	250

6 or 7 and never incur entry. With MLs choosing 6 or 7, MHs cannot profitably imitate them as 2 dominates 6 and 7 for MHs. Once again beliefs supporting these equilibria are that deviations to outputs used with zero probability in equilibrium involve an MH type with sufficiently high probability to induce entry. This deters MLs from choosing lower output levels. *These separating equilibria involve strategic play (limit pricing) by MLs since output levels 6 and 7 are higher than would be ideal if the effect on E's response is ignored.*

For the limit pricing game with low cost Es (Tables 1, 3) the expected value of IN is greater than OUT if both types choose the same output level. This destroys any pure strategy pooling equilibrium, leaving the two pure strategy separating equilibria just described. Also playing a role in the experimental data is a mixed strategy equilibrium where MHs choose 2 with probability 0.80 and 5 with probability 0.20 and MLs always choose 5. *This too involves strategic play by MLs as they choose a higher output level than would be optimal ignoring E's response.*

3 Experimental procedures and design

We begin this section by describing the general procedures used in all sessions and then lay out the specifics of the experimental design.

3.1 General procedures

Experimental sessions employed between 12 and 16 subjects who were randomly assigned to computer terminals. All sessions included an even number of subjects so all individuals could play in every round. For inexperienced subject sessions, a common set of instructions were read out loud, with each subject having a written copy. Subjects had copies of both Ms' and Es' payoff tables and were required to fill out short questionnaires to insure their ability to read them. After reading the instructions, questions were answered out loud and play began with a single practice round followed by more questions. At the beginning of experienced subject sessions, an abbreviated version of the full instructions was read out loud with each subject having a written copy.⁵

Before each play of the game the computer randomly determined each M's type and displayed this information on Ms' screens. Ms chose first, with each M's choice sent to the E they were paired with for that game. Es then decided between IN and OUT. Following each play of the game subjects learned their payoffs and Es were told the type of M they were paired with. In addition, the lower left-hand

⁵ A copy of the instructions is available at <http://www.econ.ohio-state.edu/kagel/>

portion of each subject's screen displayed the results for each pairing: M's type, M's action, and E's response. Thus, subjects had a full history of Ms' actions conditioned on their type and Es' responses conditioned on Ms' actions. Subject ID numbers were suppressed throughout to preserve anonymity.

To speed learning, subjects switched roles after every 6 games, with Ms becoming Es and vice versa. We refer to a block of 12 games with each subject playing each role for 6 games as a "cycle." Within each set of 6 games, each M was paired with a different E for every play of the game.

All but one inexperienced subject session had 36 games, with the number of games announced in advance.⁶ Experienced subject sessions had a minimum of 36 games, with all but two of the control sessions having 48 games. All of the crossover sessions use experienced subjects.

When a crossover took place all subjects were given written copies of the new payoff tables. A brief set of instructions were read out loud indicating that the basic structure of the game was the same as before but that payoffs had changed. The number of additional games to be played was also announced.

Subjects were recruited through announcements in undergraduate classes, posters placed throughout the University of Pittsburgh and Carnegie Mellon University, and advertisements in campus newspapers. This resulted in recruiting a broad cross section of undergraduate and graduate students from both campuses. Sessions lasted a little under two hours. Subjects were paid \$5 for showing up on time. Earnings averaged \$17.50 per subject in inexperienced subject sessions. Earnings were generally higher in experienced subject sessions, largely as a result of playing more games.⁷

At the end of the inexperienced sessions, subjects were asked if they were interested in returning for a second session. Experienced subject sessions generally took place about a week after the inexperienced subject sessions. Subjects from different inexperienced subject sessions were mixed in the experienced subject sessions.

Control sessions were conducted using both a "generic" context and a "meaningful" context. The generic context uses abstract terms throughout. For example, monopolists are referred to as "A players," with the two types being "A1 types" and "A2 types" respectively, and potential entrants are described as "B players." Other terms are given similarly meaningless labels. The meaningful context uses natural terms while avoiding any value laden language. Thus, the monopolist is referred to as the "existing firm," with the two types being "high cost firms" and "low cost firms" respectively, and the potential entrant becomes the "other firm" deciding between entering "this" market or some "other" market. No subject was ever switched between generic context and meaningful context or vice versa. All crossover sessions used meaningful context. In an earlier paper we find that meaningful context speeds up learning for inexperienced subjects (Cooper and Kagel

⁶ One session had only 24 games since it was conducted in an undergraduate economics class during class time, which limited the number of games. Subjects from this session were slightly *more* likely than other experienced subjects in the control sessions to play strategically as MLs (76 vs. 61%).

⁷ These sessions also tended to be shorter since only an abbreviated version of the instructions were read and subjects were familiar with the game.

2003), but does not affect the play of experienced subjects. We control for any potential context effects in the statistical analysis.⁸

3.2 Experimental design

Our experimental design compares the development of strategic play by MLs between “crossover” sessions where Es’ payoffs switch from Table 2 (high cost Es) to Table 3 (low cost Es) versus control sessions where Es use Table 3 (low cost Es) throughout.

There were three crossover sessions with a total of 38 subjects. All of these subjects had participated in at least one full session of the limit pricing game with high cost Es. One prediction of the fictitious play model (without sophisticated learners) developed in Sect. 5 is that the frequency of strategic play by MLs following the crossover is a decreasing function of experience in the game with high cost Es. To test this prediction, the crossover to the game with low cost Es occurred at different times. In one session subjects were crossed in the 13th game after returning as experienced subjects. In a second session, all subjects played in a full *experienced* subject session with high cost Es before playing in a third session in which they were crossed to the low cost E game in the 13th game.⁹ In the third session half the subjects had played one prior session with high cost Es and half had played two prior sessions with high cost Es. This session was crossed to low cost Es in the 25th game.

There were 5 experienced subject control sessions with a total of 66 subjects. Only subjects who returned for an experienced subject control session are included in the data set for inexperienced subject controls sessions (to avoid comparing subjects who returned with those who did not).

Past experiments with the limit pricing game with high cost Es find that play reliably converges to the pooling equilibrium at output 4 (Cooper et al. 1997b). Strategic play in this game involves MHs imitating MLs by choosing output levels 3, 4, or 5. Introducing low cost Es (Tables 1, 3) eliminates all pooling equilibria. Strategic play now requires MLs to choose output levels 5, 6, or 7, distinguishing themselves from MHs. While the actions used to play strategically are changed following the crossover, the concepts underlying strategic play in the game with high cost Es remain relevant. To play strategically in either game, Ms must realize that the Es will be trying to infer their types from their output choices and that by choosing a relatively high output they can make themselves seem more like a low cost type.

4 Results

Our experimental design relies on the emergence of a pooling equilibrium in the limit pricing game with high cost Es prior to the crossover. In the last twelve period cycle before Es’ payoffs changed, 60.2% of play by MHs is at 4, and 67.3% involves

⁸ The context controls allow us to identify that there is positive transfer in Experiment 2 even if the crossover sessions are directly compared only to the control sessions employing meaningful context. Cooper and Kagel (2006) studies the interaction between using meaningful context and the crossover effect.

⁹ One subject was once-experienced in this crossover session. She was needed to make an even number of players.

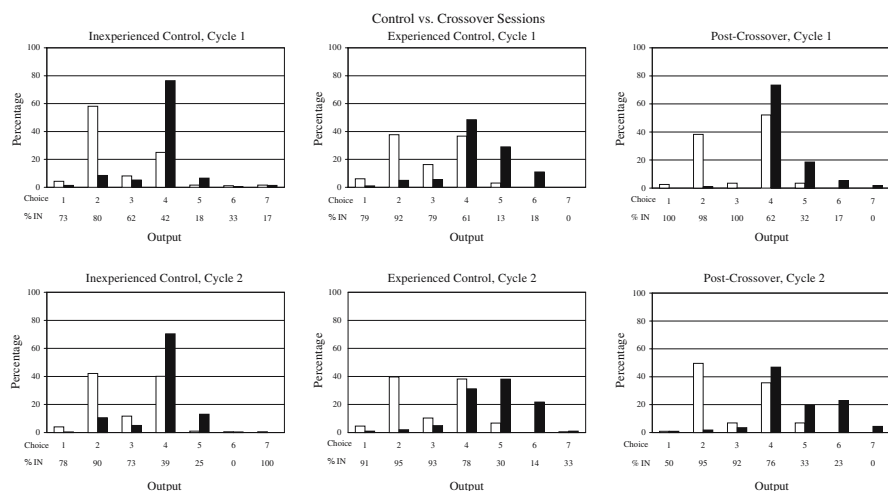


Fig. 1 Control vs. crossover sessions

strategic play of some sort (choice of 3, 4, or 5). For MLs, 89.6% of play is at 4 and there are almost no choices at higher output levels (5.2%). These choices are supported by strong incentives to limit price as an MH but not as an ML, as the entry rate differential between 2 and 4 had risen to 72.4%, while the entry rate on 4 had fallen to 7.6% at most. In both cases a 13% entry rate differential is needed to support strategic play.¹⁰

The left and center panels of Fig. 1 illustrate the development of strategic play by MLs in the control sessions. The left panels show the distribution of choices for inexperienced subjects in the control sessions. Initially, MLs overwhelmingly choose the myopic maximum, output level 4. Not only is there little strategic play by MLs, it is difficult to eliminate pure errors as a cause of this strategic play since MLs' choice of output levels below 4 are more frequent than strategic play in the first cycle of play. Comparing the first and second cycles of play in the inexperienced control sessions, strategic play by MLs increases but at a slow pace. Choice of output level 4 remains the modal choice for MLs by a wide margin and choice of output levels below 4 continues to be almost as common as strategic play. This failure of MLs to play strategically cannot be attributed to a lack of incentives, as the expected payoffs for output levels 5 and 6 are both higher for MLs than the expected payoff from output level 4 (MLs' expected payoffs are 520, 541, and 592 for output levels 4, 5, and 6, respectively). The middle panels of Fig. 1 show the distribution of choices in the first two cycles of the experienced control sessions. Strategic play by MLs continues its slow growth, fueled by increasing incentives to behave strategically. Only in the second cycle does output level 4 cease to be the modal choice of MLs. To summarize:

Conclusion 1: play in control sessions starts with Ms largely choosing their respective myopic maxima, with strategic play by MLs (play of output levels 5–7) developing only gradually.

¹⁰ There are no choices of 6 in this cycle, so we cannot calculate the entry rate differential between 4 and 6. However, it cannot be greater than the 7.6% entry rate for 4.

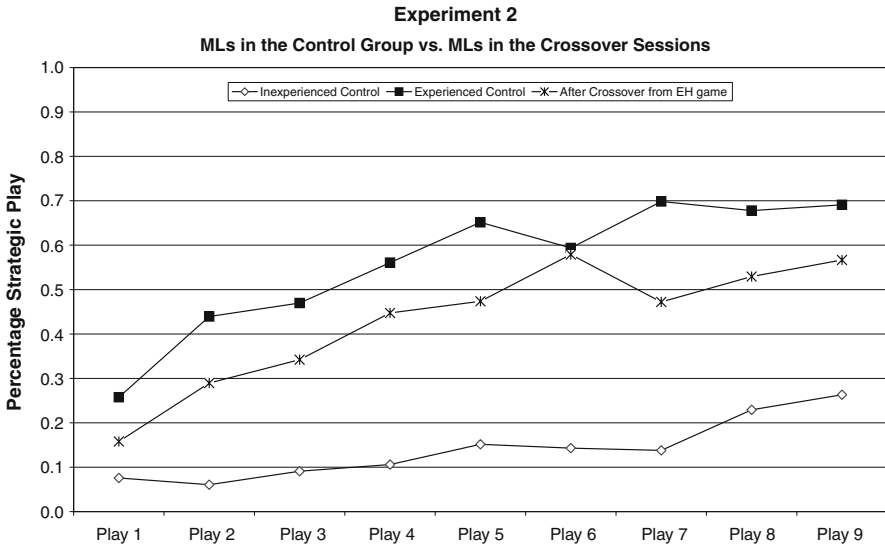


Fig. 2 Experiment 2: MLs in the control group vs. MLs in the crossover sessions

The right panels of Fig. 1 show the distribution of choices in the first two cycles following the crossover. In the first cycle of play following the crossover, output level 4 remains the modal choice for MLs. However, strategic play by MLs is three times more frequent than in the first cycle of inexperienced subject play for the control group (25.7 vs. 8.5%). Indeed, inexperienced subjects in the control group never achieve the level of strategic play observed for MLs in the first cycle following the crossover. It is only in the first cycle of *experienced* subject play for the control group that we see more strategic play by MLs (40%) than immediately following the crossover.

Figure 2 provides a more detailed view of the crossover effect for MLs. The unit of time on the x-axis is how many times a subject has played as an ML. On average, each subject will have three such plays in a twelve period cycle. In control sessions, time is measured from the beginning of the session. For example, “Play 1” is the first time an inexperienced subject played as an ML. In crossover sessions time is measured from the point of the crossover. In this case “Play 1” is the first time a subject played as an ML following the crossover. The graph plots the percentage of strategic play by MLs in inexperienced control sessions, experienced control sessions, and crossover sessions. Looking at “Play 1,” MLs in the crossover treatment immediately limit price more often than their counterparts in the inexperienced control sessions.¹¹ This suggests that MLs anticipate a change in Es’ behavior following the crossover. The evolution of play by MLs in the crossover treatment closely parallels that of the experienced control group, diverging steadily from the inexperienced control group. These results can be summarized as follows:

¹¹ It doesn’t matter much how long subjects wait before their first opportunity to play strategically as an ML. A simple way to see this is to compare the behavior of subjects who are initially Ms following the crossover with those who are initially Es – the latter uniformly wait longer to play as MLs. The percentage of strategic play is identical across these two groups in their first opportunity as an ML.

Conclusion 2: MLs in the crossover treatment look more like experienced than inexperienced subjects in the control sessions. Experience in a game with high cost Es appears to be a good substitute for experience in a game with low cost Es.

A confounding factor here is the greater incentives for MLs to limit price following the crossover than in the control sessions. However, as the formal statistical analysis in the appendix shows, MLs' higher frequency of strategic play following the crossover remains statistically significant after controlling for entry rate differences. The statistical analysis also addresses another potential confound, the use of both generic and meaningful context in the control sessions. The positive crossover effect is still significant with the addition of context controls.¹²

The formal statistical analysis addresses several other questions. First, it shows that the timing of the crossover does not have a significant effect on the frequency of strategic play following the crossover. This is evidence against fictitious play learning, absent sophisticated learners. Second, having been paired with an ML who played strategically prior to the crossover has no significant effect on the frequency of strategic play following the crossover. This rules out imitation as explaining the jump in strategic play following the crossover. However, playing strategically prior to the crossover as *either* an ML *or* an MH is positively and significantly correlated with strategic play as an ML following the crossover. Although this result in part reflects individual effects in the data, it also draws on deeper aspects of subjects' learning processes. The structural model of learning developed in Sect. 5 indicates that rapid development of strategic play following the crossover is closely tied to the presence of "sophisticated learners" in the population. Sophisticated learners are more likely to play strategically for the game with high cost Es and the game with low cost Es, thereby helping to generate positive correlation between individual subject's strategic play before and after the crossover.

5 A structural model of learning and sophistication

The positive cross-game learning reported above is inconsistent with the predictions of the stochastic fictitious play learning model that motivated the experiment. This model predicts negative cross-game learning because Ms fail to anticipate the change in Es' behavior following the crossover and therefore respond incorrectly. In contrast, the experimental results suggest that *at least some* subjects are sophisticated enough to anticipate the change in Es' behavior and correctly respond to it.

To explore this intuition more formally, this section develops a stochastic fictitious play learning model, adds sophisticated learners, and fits both models to the data. This analysis has two purposes. First, we want to show that the addition of sophisticated learners improves the econometric fit to the data. Second, and more importantly, we want to show that the addition of sophisticated learners enables the learning model to track the main features of MLs' behavior following the crossover.

The basic learning model treats Ms as belief-based learners in the spirit of stochastic fictitious play (Fudenberg and Levine 1995). We choose this model because

¹² Additionally, see Cooper and Kagel (2006) for a replication of the positive crossover effect using only meaningful context for control and crossover sessions.

similar models have done a good job of tracking the development of play in earlier signaling game experiments (Cooper et al. 1997a,b). The model, although only described for Ms in our limit pricing game, generalizes in a straight forward way to other games.

We have not explicitly considered other classes of learning models such as replicator dynamics (see Fudenberg and Levine 1998, chap. 3), reinforcement learning (Roth and Erev 1995), or EWA (Camerer and Ho 1999). Determining the learning model that best tracks subjects' behavior goes well beyond the scope of the present paper. It is unlikely that using these other models would overturn our main conclusions since all of them, like fictitious play, embody unsophisticated learners who do not explicitly model other players' learning and decision making processes.

5.1 The Learning Model

A belief-based learning model requires rules for choosing a strategy in period t given beliefs, updating beliefs from period t to period $t + 1$, and generating initial beliefs. Let $C_{ij}^t(\text{IN})$ and $C_{ij}^t(\text{OUT})$ be weights that player i puts on the responses "IN" and "OUT", respectively, in period t following output j . These variables can be thought of as modified counts for the number of times each outcome has been observed. Let $b_{ij}^t(\text{IN})$ and $b_{ij}^t(\text{OUT})$ be the probabilities that player i assigns to the responses "IN" and "OUT" respectively in period t following output j . These represent player i 's beliefs. Beliefs are generated from $C_{ij}^0(\text{IN})$ and $C_{ij}^0(\text{OUT})$ using the following two equations:

$$b_{ij}^t(\text{IN}) = \frac{C_{ij}^t(\text{IN})}{C_{ij}^t(\text{IN}) + C_{ij}^t(\text{OUT})} \quad (1a)$$

$$b_{ij}^t(\text{OUT}) = 1 - b_{ij}^t(\text{IN}) \quad (1b)$$

Given $b_{ij}^t(\text{IN})$, $b_{ij}^t(\text{OUT})$, and player i 's type in period t , let π_{ij}^t be player i 's expected payoff from choosing output j in period t . With probability p_{change} , player i selects a new strategy in period t . Otherwise, he uses the same output as the last time he played as the same type.¹³ Player i 's probability of choosing output j in period t (subject to choosing a new output), p_{ij}^t , is generated via a logit rule:

$$p_{ij}^t = \frac{e^{\lambda \pi_{ij}^t}}{\sum_{k=1}^7 e^{\lambda \pi_{ik}^t}} \quad (2)$$

This rule has the usual interpretation. The precision parameter λ is the level of noise in the system. If $\lambda = 0$, the result is pure noise with each strategy chosen with equal probability. As $\lambda \rightarrow \infty$, we get arbitrarily close to best-response to beliefs.

Individuals learn by updating $C_{ij}^t(\text{IN})$ and $C_{ij}^t(\text{OUT})$ from period to period. Some notation is required before the updating rule can be written down. Let δ be

¹³ In previous fitting exercises (Cooper and Stockman 2002; Stahl 2003), introducing autocorrelation into the model significantly improved the fit.

the discount rate for past experience. Define $C_{ij}^t(\text{IN})$ and $C_{ij}^t(\text{OUT})$ to be the number of times that player i chose output j in period t and observed the responses “IN” or “OUT” respectively. Define $c_{-ij}^t(\text{IN})$ and $c_{-ij}^t(\text{OUT})$ to be the number of times that an M player other than player i chose output j in period t and observed the responses “IN” or “OUT”, respectively. Finally, given that subjects see the results for all other pairings, let w_{Other} be the weight players put on the experience of other players relative to their own experiences. The updating rule for $C_{ij}^t(\text{IN})$ in periods with no crossover is given by Eq. 3, with the updating rule for $C_{ij}^t(\text{OUT})$ defined in an analogous manner. Note that updating takes place even in periods where player i is not playing as an M.

$$C_{ij}^{t+1}(\text{IN}) = \frac{C_{ij}^t(\text{IN})}{1 + \delta} + c_{ij}^t(\text{IN}) + w_{\text{Other}} \cdot c_{-ij}^t(\text{IN}) \quad (3)$$

For periods following a crossover, the updating rule accounts for the possibility that subjects will “reset” their beliefs. In other words, beliefs following the crossover are treated as a convex combination of beliefs prior to the crossover and the beliefs of an inexperienced subject. Suppose a crossover takes place between period t and period $t + 1$. Let ρ be the weight on resetting beliefs. Player i ’s beliefs are first updated using (3). The following additional transformation is then made, where $C_{ij}^t(\text{IN})$ gives the counts prior to the transformation and $C_{ij}^t(\text{IN})$ gives the counts following the transformation. An analogous transformation is made for $C_{ij}^t(\text{OUT})$.

$$C_{ij}^t(\text{IN}) = (1 - \rho)C_{ij}^t(\text{IN}) + \rho C_{ij}^0(\text{IN}) \quad (4)$$

Intuitively, $(1 - \rho)$ gives the weight subjects put on experience from the previous related game. If $\rho = 1$ there is no cross-game learning and if $\rho = 0$ the games are treated as being identical.¹⁴

To generate initial values for $C_{ij}^0(\text{IN})$ and $C_{ij}^0(\text{OUT})$, we fit initial beliefs for each of the seven strategies. Since probabilities must add up to 1, this involves fitting seven parameters. We then fit a single variable, “Strength,” that determines the initial strength of beliefs for all seven strategies. $C_{ij}^0(\text{IN})$ and $C_{ij}^0(\text{OUT})$ are backed out of the fitted parameters. Let $b_j^0(\text{IN})$ be the initial belief that an E will enter following output level j . Then $C_{ij}^0(\text{IN}) = b_j^0(\text{IN}) \cdot \text{Strength}$ and $C_{ij}^0(\text{OUT}) = \text{Strength} - C_{ij}^0(\text{IN})$.

Having described the basic learning model, we now modify it to include two additional modes of learning: non-learners and sophisticated learners. Non-learners start with the same initial beliefs as unsophisticated learners, make choices in exactly the same way as unsophisticated learners, but never update their beliefs. A sophisticated learner models Es as being unsophisticated learners who maximize payoffs subject to their beliefs. This implies that a sophisticated learner anticipates that changes in payoffs will affect Es’ choices, and that Es’ behavior will change as they accumulate experience.

In choosing how to incorporate sophistication into the learning model, our goal is to use the minimal level of sophistication necessary to track the data. The type

¹⁴ We have explored a variety of other specifications for how beliefs might be transformed following the crossover. The qualitative results are unaffected by alternative specifications.

of sophistication we have added represents a relatively modest change to the stochastic fictitious play model. This approach has a number of antecedents in the literature, particularly Milgrom and Roberts (1991), Selten (1991), Nagel (1995), Stahl (1996) and Camerer et al. (2002). Its key role is to allow Ms to anticipate changes in entry rates following the crossover. Although subjects may in fact be operating at a higher level of sophistication, learning very general concepts about signaling games, the data does not force us to this conclusion. The level of sophistication added to the model does not imply that subjects can generalize what they have learned in the limit pricing game to a radically different signaling game any better than inexperienced subjects. For example, our sophisticated learners would not necessarily be able to perform any better than inexperienced subjects in Brandts and Holt's (1992) signaling game or in Miller and Plott's (1985) game as they have substantially different structures from the present game. It remains an open empirical question whether or not the higher levels of sophistication needed for such cross-game learning exist in the population.

Going into the details, a sophisticated learner needs to build beliefs that best replicate the beliefs an unsophisticated E might have. These are not the sophisticated player's beliefs, rather they are his best estimate of an unsophisticated E's beliefs. He builds these beliefs in exactly the same manner that an unsophisticated E would. In estimating the beliefs of unsophisticated Es, updating is done in a manner analogous to (3) and (4) above, but with one important difference—outcomes from other players are weighted equally to a player's own outcomes. Intuitively, a sophisticated learner is building fictitious beliefs for other players and therefore has no reason to overweight his own experience. Given his best estimate of the beliefs of Es, a sophisticated learner generates a probability of entry for each output level using a logit rule analogous to (2). The resulting probabilities give a sophisticated learner's beliefs about the behavior of Es. Based on these beliefs, a sophisticated learner generates his own choice in exactly the same manner as an unsophisticated learner. Thus, a sophisticated learner uses a noisy best response to a noisy best response to beliefs based on observed outcomes.¹⁵

The model allows players to switch modes of learning over the course of play. To simplify computations, the only time this switch is allowed is when players return as experienced subjects. We further simplify the model by only allowing players to move up a single level of sophistication. Thus, there are three "pure" types (non-learners, unsophisticated learners, and sophisticated learners) and two "switching" types (non-learner to unsophisticated learner and unsophisticated to sophisticated learner). The *ex ante* probabilities of these five types are parameters that we fit from the data.¹⁶

¹⁵ The model can be modified to allow for types who anticipate a mixture of other types or use a mixture of sophisticated and unsophisticated learning. While this would no doubt improve the model's ability to fit the data, it complicates the model while adding little to our understanding of the underlying cognitive processes.

¹⁶ The fitting exercise does not assign specific types to the subjects. Instead it generates the likelihood of a subject's observed choices subject to being a certain type, and then generates the full likelihood by taking the weighted average over types, where the weights are given by the *ex ante* probability of each type. Allowing players to switch types at more points in time generates a statistically significant improvement in the fit, reflecting the presumably continuous nature of switching in reality, but does not change the qualitative results. Allowing types that switch up more than one level of sophistication or types that switch to *lower* levels of sophistication does not generate a statistically significant improvement in the fit.

5.2 Fitting the model

The model was fit using data from all subjects (including the controls) who returned for an experienced subject session. All plays as an M in both roles are used. Parameter estimates are generated through standard maximum likelihood techniques, with probabilities bounded between 0 and 1. When the algorithm ran into the boundaries for parameters that represent probabilities, they were set equal to the boundary values to allow for convergence.

We set the initial beliefs for subjects playing in games with high cost Es and those for subjects playing in games with low cost Es equal, as a log likelihood ratio test fails to reject the null hypothesis of identical initial beliefs ($\chi^2 = 9.68$, $df = 7$, $p > 0.10$). To simplify computations, the following parameters are set equal (where relevant) for all three behavioral types: the precision parameter (λ), the probability of changing strategies (p_{change}), discounting of past experience (δ), and the reset parameter (ρ). In addition, the initial beliefs sophisticated learners assign to unsophisticated Es are forced to be identical across low outputs (1 and 2), intermediate outputs (3 and 4), and high outputs (5, 6, and 7). Relaxing these restrictions would strengthen our main conclusions, but makes the likelihood function substantially harder to maximize.

The data set includes repeated observations from the same individuals which cannot be treated as statistically independent. The inclusion of “inertia” in the model through the variable p_{change} somewhat controls for these individual effects. The inertia variable adds correlation between observations from the same individual, so that its effect is roughly analogous to what a random effect specification does in more standard sorts of analysis. To the extent that this does not account for all of the individual effects in the data, we also apply the correction for clustering suggested by Moulton (1986) and White (1994) to the standard errors.

The results of the maximum likelihood estimation are reported in Table 4. Standard errors (corrected for clustering) are shown in parentheses. The estimates of initial beliefs are suppressed since these are of little direct interest. Results from three versions of the model are reported. Model 1 only includes non-learners. Model 2 includes non-learners and unsophisticated learners, with no switching between non-learners and unsophisticated learners. When a probability of switching is included in Model 2, the maximization algorithm sets it equal to zero (indicating that it can be deleted). Model 2 is equivalent to a standard model of stochastic fictitious play. Model 3 is the full model with non-learners, unsophisticated learners, and sophisticated learners, as well as switching between types.

Comparing Model 1 with Model 2, we see a large improvement in the log-likelihood ($\chi^2 = 603.60$, $df = 5$, $p < 0.01$). Not surprisingly, given the strong dynamics in the data, the evidence in favor of learning is overwhelming. The improvement in the log-likelihood between Model 2 and Model 3 is also large and significant at the 1% level ($\chi^2 = 189.42$, $df = 6$, $p < 0.01$). Looking at the parameter estimates, the proportion of sophisticated learners increases from 18.8% in the inexperienced sessions to 32.4% in the experienced sessions, and the estimated proportion of non-learners falls from 25.4% in the inexperienced sessions to 18.1% in the experienced sessions. Even though the latter decrease is not

Table 4 MLE Results for learning models standard errors corrected for clustering

	Model 1	Model 2	Model 3
Properties of the model			
Non-learners	✓	✓	✓
Unsophisticated learners		✓	✓
Sophisticated learners			✓
Switching between types			✓
Number of parameters	9	14	20
Parameter estimates			
Precision (λ) (Multiplied by 100)	1.534** (0.166)	1.960** (0.111)	2.384** (0.117)
Probability change of strategy (p_{change})	0.503** (0.023)	0.645** (0.029)	0.674** (0.026)
Discounting of past experience (δ)		0.084** (0.009)	0.060** (0.010)
Square root of weight on initial beliefs following crossover (ρ)		0.054+ (0.028)	0.022 (0.036)
Weight on others' experience (w_{Other}) (Multiplied by 100)		0.259 (0.162)	0.699 (0.454)
Probability non-learner		0.310** (0.053)	0.254** (0.046)
Probability sophisticated learner			0.188** (0.044)
Probability switching type non-learner to unsophisticated			0.073 (0.072)
Probability switching type unsophisticated to sophisticated			0.136** (0.045)
Log likelihood	-4517.11	-4215.31	-4120.60

The full data set has 4,595 observations over 104 individuals, including 2,585 observations from 66 individuals in the control sessions and 2,010 observations from 38 individuals in the crossover sessions

- ** Statistically significant at the 1% level
- * Statistically significant at the 5% level
- + Statistically significant at the 10% level

statistically significant, the population is clearly moving toward greater sophistication over time.¹⁷

Conclusion 3: The addition of “sophisticated” learners to the basic model of fictitious play generates a statistically significant improvement in the fit to the data.

5.3 Simulations

This subsection reports simulations showing that the learning model without sophisticated learners misses important features of the data that the model with sophisticated learners captures. Thus, the addition of sophisticated learners to the model is not just statistically significant, it is economically significant as well.

We simulate Ms’ learning using the parameters generated by the maximum likelihood estimation. The simulations are designed to closely mimic the experiment. Since we are primarily interested in the strategic play of MLs, the responses

¹⁷ The estimate for the “reset” parameter (ρ) in Model 3 is small and not significant indicating that unsophisticated learners’ beliefs are almost unaffected by the crossover. The estimate of “weight on others’ experience” is small and statistically insignificant suggesting that unsophisticated learners’ beliefs are based primarily on their own experience. The probability of changing strategies is always significantly less than 1, implying autocorrelation.

of Es and MHs are generated randomly using the observed frequencies in the data. Simulations were run for inexperienced subject sessions with 36 games and experienced subject sessions with 48 games, with the crossovers taking place in game 13. As in the experiment, simulated subjects alternated between playing as Ms and Es, with half of the simulated subjects as Ms for the first half of each twelve period cycle and the other half as Ms in the second half. One slight difference from the experiment is that we forced each simulated player to be an ML (MH) exactly three times in each twelve period cycle. For each model and each treatment, play was generated for 10,000 simulated subjects for each of the five behavioral types (including the two switching types). The fitted probabilities of each type were then used to generate aggregate behavior.

Figure 3 displays strategic play by MLs from the simulations in the same way that Fig. 2 did for the experimental data. The unit of time on the x -axis is how many times a subject has played as an ML. The top panel reproduces the data from the experiment (Fig. 2), the middle panel simulates play without sophisticated learners, and the bottom panel simulates play with sophisticated learners.

Comparing the top and middle panels of Fig. 3, the simulated subjects do not replicate the immediate jump in strategic play by MLs that is observed in the data following the crossover. Intuitively, unsophisticated learners have no mechanism to quickly adjust their beliefs about Es' behavior following the change in their payoffs. The only way the model without sophisticated learners can even partially replicate MLs' rapid jump to strategic play following the crossover is by allowing for very fast learning following the crossover. (This is the reason we estimate higher values of the discount parameter δ and the reset parameter ρ in the model without sophisticated learners than with sophisticated learners.) These simulations confirm that, even fitted to data from the crossover sessions, the learning model without sophisticated learners cannot track the data.

In contrast, the simulations with sophisticated learners look similar to the experimental data: simulated MLs immediately show more strategic play following the crossover than simulated inexperienced subjects in the control treatment. Subsequently, strategic play by simulated MLs in the crossover treatment grows gradually, paralleling the growth of strategic play for experienced subjects in the control treatment. Thus, the addition of sophisticated learners not only improves the statistical fit to the data, it allows us to track the major features of play following the crossover.

The presence of sophisticated learners who immediately anticipate the effect of changing Es' payoffs is necessary but not sufficient to explain why there is significantly *more* strategic play by MLs following the crossovers than in the inexperienced control sessions. Without the increase in sophisticated learners following the crossover, we would see almost exactly the same level of strategic play as in the inexperienced control sessions.¹⁸ In the learning model, experience with the game with high cost Es helps generate more strategic play by MLs *because the level of sophistication has grown over time in the subject population as a result of playing a*

¹⁸ We reran the simulations of Model 3 not allowing for any growth in the proportion of sophisticated learners. Comparing the first play as an ML in the inexperienced control sessions with the first play as an ML following the crossover, we see only a 2% increase in the frequency of strategic play. This is far smaller than the 8% difference observed in the actual data or the 12% difference in the simulations with switching between types.

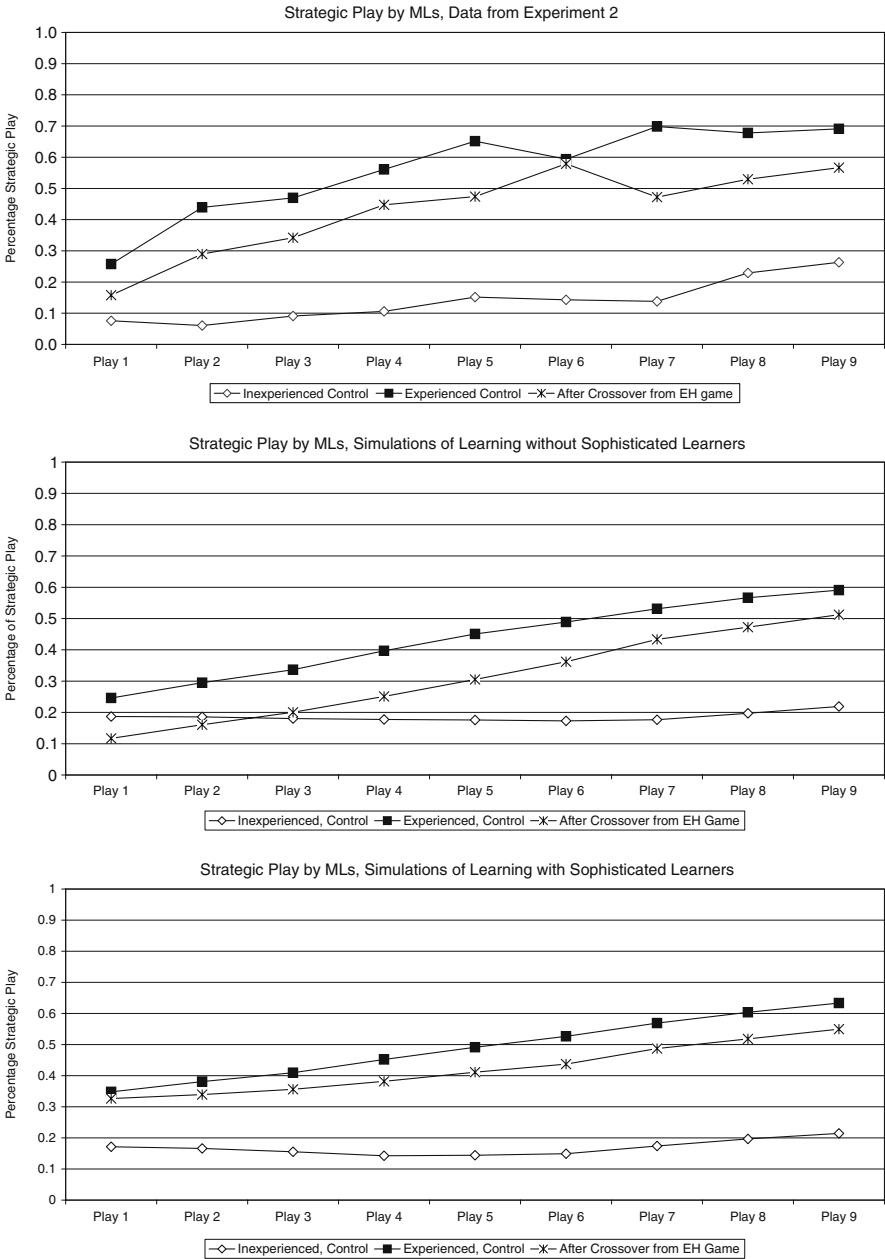


Fig. 3 Comparison of actual and simulated play by MLs

related game. That is, the critical difference between the crossover sessions and the inexperienced control sessions is that experience with the high cost entrant game results in a higher percentage of sophisticated learners in the population. Thus, the

primary mechanism underlying the surprising degree of positive transfer following the crossover is the *growth* in sophistication in the subject population.

Conclusion 4: The learning model with sophisticated learners generates better tracking of MLs' behavior in crossover sessions. Growth over time in the proportion of sophisticated learners provides a mechanism for the positive transfer observed in crossover sessions.

6 Summary and conclusions

This paper studies cross game learning in signaling games. Study of cross game learning is important since, as Fudenberg and Kreps (1988) note:

“... it seems unreasonable to expect the exact same game to be repeated over and over; put another way, if we could only justify the use of Nash analysis in such situations, we would not have provided much reason to have faith in the widespread applications that are found in the literature. Faith can be greater if, as seems reasonable, players infer about how their opponents will act in one situation from how opponents acted in other, similar situations.”

Our experiments provide evidence that subjects who have learned to play strategically in one game can transfer much of this knowledge to related games even if the actions necessary to play strategically are quite different. More importantly, we have begun to understand the mechanism(s) underlying this transfer. We find evidence that there exist sophisticated learners in the subject population and that the proportion of sophisticated learners increases with experience. This growth in sophistication plays a central role in fostering transfer. In other words, experience not only changes how subjects play games, but also how they approach related games, generating increased sensitivity to the strategic implications of their actions and the effects of changes in other player's payoffs. It is this increased sensitivity that allows them to perform well compared to naive subjects when put into a new (but related) setting.

The dynamic mechanism underlying the learning transfer explains why the negative results in the psychology literature aren't replicated here. The psychology literature focuses on one shot trials where subjects learn specific skills (e.g. how to drive a truck) or how to solve certain classes of problems (e.g. logic puzzles). Games, by their nature, are interactive. With experience, subjects gain the ability to think about how other individuals are making decisions and incorporate this into their own decision making. Both the interactive element of games and the extended experience necessary to generate sophisticated reasoning about games are missing from the individual choice problems studied by psychologists.

We attribute the positive transfer following the crossover to the existence of a growing population of sophisticated learners, but our experimental design does not allow us to directly verify the existence (and increasing frequency) of sophisticated learners since we have no direct observations of subjects' cognitive processes. In subsequent research in which two person teams play the roles of Ms and Es, content analysis of communication between team members verifies (i) the existence of sophisticated learners of the type modeled here and (ii) growing numbers of sophisticated learners as a result of experience with related games (Cooper and Kagel 2005).

Finally, ongoing research shows that our use of meaningful context for the cross-over sessions plays a role in the development of sophisticated reasoning and, by extension, the occurrence of positive cross-game learning (Cooper and Kagel 2006). Using generic context, the same cross-over treatment generates zero or even slightly negative cross-game learning. This positive effect of meaningful context on cross-game learning is not universal, as Cooper and Kagel (2006) also report an example where the effect is negative. In brief, these divergent results capture different aspects of the mechanism underlying positive cross-game learning. The experiments above stress one channel for generating positive cross game learning; in the initial game subjects acquire concepts underlying strategic play that are applicable to the subsequent game. Meaningful context helps in this process by fostering the development of strategic empathy. However, if the main barrier to cross-game learning is the ability to recognize that experiences in the initial game are relevant for the new game, meaningful transfer may play a negative role by obscuring the relationship between games.

Appendix

This appendix shows that Conclusion 2 is supported by formal econometric analysis of the data controlling for covariates affecting behavior.

The regressions reported in this appendix are ordered probits. The use of an ordered probit specification recognizes that the output choices by Ms are inherently categorical data. There are two reasons for this. First, suppose that the subjects have preferences over a continuum of possible output choices. Because the design only allows them seven possible choices, individuals whose true preferences differ may end up in the same category. For example, suppose that one subject most preferred output level is 4.8 and another's is 5.2. These may both show up in the data as a choice of output level 5. The use of an ordered probit explicitly accounts for the mapping between a discreet choice set and an underlying continuous space of possible choices. Second, the game itself is fundamentally non-linear. For example, consider the difference as an ML between moving from output level 5 to 6 and moving from 6 to 7. Beyond any strategic considerations, just considering the payoffs, the later is a much larger change than the former. The non-linearity of an ordered probit captures the idea that not all changes of a single output level are equal.

The dependent variable in all of the regressions is the output level chosen by MLs. To correct for individual effects in the data, standard errors are calculated using Moulton's (1986) correction for clustering. In addition to the ordered probits reported here, we have run a variety of other specifications including linear models with a random effects specification, probits with a random effects specification, and ordered probits with a limited number of categories and a random effects specification. Our qualitative conclusions are the same for any of these alternative approaches to the data.

Table 5 reports the regression results. The data set for these regressions includes all data from games with low cost Es for subjects who returned for an experienced subject session. Data from games with high cost Es (data prior to the crossover) are not included.

Table 5 Crossover effects for MLs: ordered probits, standard errors corrected for clustering, dependent variable is output level

Control variables	Model 1	Model 2	Model 3	Model 4
	None	Entry rates	Entry rates context	None
Number of parameters	15	20	27	19
(β_3) No Crossover, experienced periods 1–12	0.613** (0.151)	0.750** (0.213)	0.395+ (0.226)	0.640** (0.158)
(β_4) No Crossover, experienced periods 13–24	1.095** (0.155)	1.268** (0.198)	0.902** (0.197)	1.145** (0.163)
(β_5) No Crossover, experienced periods 25–36	1.323** (0.161)	1.418** (0.176)	1.053** (0.201)	1.384** (0.170)
(β_6) No Crossover, experienced periods 37–48	1.534** (0.191)	1.560** (0.213)	1.291** (0.289)	1.604** (0.204)
(δ_1) Crossover periods 1–12 after crossover	0.585** (0.151)	0.592** (0.167)	0.436* (0.175)	0.529** (0.194)
(δ_2) Crossover periods 13–24 after crossover	1.052** (0.240)	1.055** (0.241)	0.915** (0.253)	1.010** (0.200)
(δ_3) Crossover periods 25–36 after crossover	1.277** (0.206)	1.355** (0.212)	1.238** (0.222)	1.380** (0.256)
(λ_1) Strategic play as MH 10 plays prior to crossover				0.123** (0.025)
(λ_2) Strategic play as ML 10 plays prior to crossover				0.235** (0.056)
(τ) Cycle when crossover occurs				0.022 (0.089)
(t) Paired with an ML who played strategically				0.238 (0.193)
Log likelihood	–1999.67	–1993.31	–1964.63	–1929.74

All regressions contain 1,654 observations over 104 individuals

** Statistically significant at the 1% level

* Statistically significant at the 5% level

+ Statistically significant at the 10% level

Formally, the full specification for the latent variable underlying the ordered probit is given by Eq. A.1 below. The variable *Lo* is a dummy for subjects who play the low cost entrant game in all periods. This is the control group. The variable *Hi* is a dummy for subjects that play the high cost entrant game initially and then are crossed over to the low cost entrant game. In other words, this is a dummy for the crossover treatment. The variables *Cyc_{I2}* and *Cyc_{I3}* are dummies for the second and third twelve period cycles of the inexperienced sessions. The variables *Cyc_{E1}*, *Cyc_{E2}*, *Cyc_{E3}*, and *Cyc_{E4}* are dummies for the first, second, third, and fourth twelve period cycles of the experienced sessions. The variables *Cyc_{CR1}*, *Cyc_{CR2}*, and *Cyc_{CR3}* are dummies for the first, second, and third twelve period cycles following the crossover from the high cost entrant game to the low cost entrant game. Thus, for subjects in the crossover treatment we control for time since the subject started playing the low cost entrant game, not total time the subject has been playing some limit pricing game. The variable *ER* is a vector of entry rate controls and the variable *Con* is a vector of controls for the use of meaningful context. The variables *SMH* and *SML* measure a subject's use of strategic play as an *MH* and as an *ML* prior to the crossover. The variable *TCRS* measures how experienced the subject was when the crossover took place. The variables *SMH*, *SML*, and *TCRS* are all set equal to zero for subjects in the control sessions. The error term is given by ε_t^i .

$$\begin{aligned} O_t^i = & \alpha + \beta_1 \text{Lo} * \text{Cyc}_{I2} + \beta_2 \text{Lo} * \text{Cyc}_{I3} + \beta_3 \text{Lo} * \text{Cyc}_{E1} + \beta_4 \text{Lo} * \text{Cyc}_{E2} \\ & + \beta_5 \text{Lo} * \text{Cyc}_{E3} + \beta_6 \text{Lo} * \text{Cyc}_{E4} + \delta_1 \text{Hi} * \text{Cyc}_{CR1} + \delta_2 \text{Hi} * \text{Cyc}_{CR2} \\ & + \delta_3 \text{Hi} * \text{Cyc}_{CR3} + \gamma \text{ER} + \eta \text{Con} + \lambda_1 \text{SMH} + \lambda_2 \text{SML} + \tau \text{TCRS} \\ & + \iota \text{IMT} + \varepsilon_t^i \end{aligned} \quad (\text{A.1})$$

Less formally, the independent variables fall into four categories as follows:

1. *Controls for the Time Period Interacted with Treatment Dummy:* The base in this specification is the first twelve period cycle of the inexperienced sessions for subjects in the control sessions. The regressions include the following dummies: inexperienced control sessions, periods 13–24; inexperienced control sessions, periods 25–36; experienced control sessions, periods 1–12; experienced control sessions, periods 13–24; experienced control sessions, periods 25–36; experienced control sessions, periods 37–48; crossover sessions, periods 1–12 following the crossover; crossover sessions, periods 13–24 following the crossover; and crossover sessions, periods 25–36 following the crossover. Parameter estimates for the inexperienced sessions are suppressed in Table 5 since they are not of any direct interest.
2. *Controls for Es' Behavior:* The entry rates for the current twelve period cycle following outputs 2, 3, 4, 5, and 6 are included as independent variables. These entry rates are calculated over all periods in the current twelve period cycle and are calculated separately for each session. To the extent that subjects' beliefs reflect the experience that they are receiving, these five variables serve as a proxy for the unobservable beliefs. Note that the measures of entry rates only use information from the current twelve period cycle, not from previous cycles. This is done for two reasons. First, based on the fitted parameters for the learning model, there is good reason to expect that subjects' beliefs will disproportionately reflect experience from recent periods. Second, and perhaps

more importantly, our central interest is in how behavior changes following the crossover. We want to know if the change in Ms' behavior following the crossover is reflecting a change in Es' behavior. We therefore need a measure of Es' behavior that emphasizes how entry rates have changed following the crossover rather than reflecting entry rates prior to the crossover. Using only the current cycle allows our measures to strongly and rapidly reflect any changes in Es' behavior following the crossover.¹⁹ As a group, the entry rate controls are always easily significant at the 1% level. These parameter estimates are suppressed in Table 5 since they are not directly relevant. The second line of the table indicates whether these variables have been included in a model.

3. *Controls for Meaningful Context:* These include a dummy for subjects who experienced meaningful context as well as interactions between the context dummy and the time dummies. The parameter estimates for these variables are not reported in Table 5, but the second line of the table shows whether these variables were included in a model.
4. *Miscellaneous:* Model 4 includes four miscellaneous independent variables. One issue is whether strategic play prior to the crossover is a good predictor for strategic play following the crossover. We therefore calculate two measures of strategic play prior to the crossover: the number of times a subject played strategically the last ten times as an MH prior to the crossover and the number of times a subject played strategically the last ten times as an ML prior to the crossover. Both of these variables are demeaned. Another natural question is whether the timing of the crossover matters. To control for when the crossover occurs, Model 4 includes a variable that measures how many twelve period cycles of experience a subject had before being crossed over. Since no subject was crossed over without at least four cycles of prior experience, we subtract four from this variable to give it a minimum value of zero. Finally, having directly observed strategic play by others might serve as a catalyst for an ML playing strategically himself. We therefore include a dummy for whether an ML in the crossover treatment was, as an E prior to the crossover, paired with an ML who played strategically. This variable is demeaned.

Turning to the results, Model 1 looks for a crossover effect without controlling for entry rates or context. The variable of primary interest here is "Crossover: Periods 1–12 After Crossover." (δ_1 in equation A.2) This parameter captures the difference between inexperienced subjects in the first twelve period cycle of the control sessions and subjects in the first twelve period cycle following a crossover. The estimate is positive and significant at the 1% level. Both of the other crossover dummies are also statistically significant at the 1% level, with the size of the parameter estimates increasing substantially over time. If we modify the specification so the other two crossover dummies ("Crossover: Periods 13–24 After Crossover" and "Crossover: Periods 25–36 After Crossover") capture differences between the second and third cycles following the crossover and the second and third cycles of the inexperienced control sessions, the two crossover dummies remain significant

¹⁹ Identical regressions have been run using a variety of alternative entry rate controls, including ones that reflect behavior in all preceding periods rather than just the current twelve period cycle. The results of these alternative regressions are similar to what is reported here.

at the 1% level.²⁰ Thus, the regression analysis confirms that there is significantly more strategic play by MLs following the crossover than in inexperienced control sessions, both in the first twelve period cycle and throughout the session.

We can change the specification of Model 1 so that the parameter estimate for “Crossover: Periods 1–12 After Crossover” captures the difference between play in the first twelve period cycle following the crossover and first twelve period cycle of the experienced control sessions. Likewise, we can also modify the specification so the other two crossover dummies capture differences between the second and third cycles following the crossover and the second and third cycles of the *experienced* control sessions. With this specification, the parameter estimate for “Crossover: Periods 1–12 After Crossover” becomes -0.028 with a standard error of 0.180 . This is not significantly different from zero. Further, no significant differences can be found between experienced control sessions and crossover sessions in later cycles either as the dummies for “Crossover: Periods 13–24 After Crossover” and “Crossover: Periods 25–36 After Crossover” both fail to achieve significance individually²¹ and the three crossover dummies fail to be jointly significant ($\chi^2 = 0.29$, $df = 3$, $p > 0.10$). Thus, there are no significant differences in strategic play between the crossover sessions and the experienced control sessions.

Model 2 adds the controls for *Es*' behavior to Model 1, and Model 3 adds the controls for context to Model 2. These additional controls are statistically significant at the 1% level ($\chi^2 = 70.08$, $df = 12$, $p < 0.01$), but have no effect on our conclusions from Model 1.

Model 4 adds the two controls for strategic behavior prior to the crossover to Model 1, the control for when the crossover took place, and the control for, as an E, having been paired with an ML who played strategically prior to the crossover.²² While neither the timing of the crossover (τ in Eq. A.2) nor direct experience with an ML playing strategically (ι in Eq. A.2) have a statistically significant effect, the parameter estimates for both of the variables measuring strategic behavior prior to the crossover (λ_1 and λ_2 in Eq. A.2) are positive and statistically significant at the 1% level.

References

- Brandts, J., Holt, C.A.: An experimental test of equilibrium dominance in signaling games. *Am Econ Rev* **82**, 1350–1365 (1992)
- Camerer, C.F.: *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press (2003)
- Camerer, C.F., Ho, T.: Experience-weighted attraction in games. *Econometrica* **67**, 827–874 (1999)
- Camerer, C.F., Ho, T.H., Chong, J.-K.: Sophisticated Experience-Weighted Attraction Learning and Strategic Teaching in Repeated Games. *J Econ Theory* **104**, 137–188 (2002)
- Cooper, D.J., Garvin, S., Kagel, J.: Adaptive learning vs. equilibrium refinements in an entry limit pricing game. *Econ J* **107**, 553–575 (1997a)

²⁰ The parameter estimates are 1.022 and 1.012 respectively with standard errors of 0.239 and 0.229.

²¹ The parameter estimates are -0.043 and -0.046 respectively with standard errors of 0.263 and 0.224.

²² We have run Model 4 with controls for entry rates and the use of meaningful context. While the resulting specification is messier, it yields qualitatively identical results to the specification shown here. We have also used different variables to measure previous experience as an E paired with a strategic ML and get qualitatively similar results.

- Cooper, D.J., Garvin, S., Kagel, J.: Signaling and adaptive learning in an entry limit pricing game. *Rand J Econ* **28**, 662–683 (1997b)
- Cooper, D.J., Kagel, J.H.: The impact of meaningful context on strategic play in signalling games. *J Econ Behav Organ* **50**, 311–37 (2003)
- Cooper, D.J., Kagel, J.H.: Are two heads better than one? Team versus individual Play in signaling games. *Am Econ Rev* **95**, 477–509 (2005)
- Cooper, D.J., Kagel, J.H.: Context and learning transfer. Mimeo (2006)
- Cooper, D.J., Stockman, C.K.: Fairness and learning: an experimental examination. *Games Econ Behav* **41**, 26–45 (2002)
- Fudenberg, D., Kreps, D.M.: A theory of learning, experimentation, and equilibrium in games. Mimeo MIT (1988)
- Fudenberg, D., Levine, D.K.: Consistency and cautious fictitious play. *J Econ Dyn Control* **19**, 1065–1090 (1995)
- Fudenberg, D., Levine, D.: *The Theory of Learning in Games*. Cambridge: MIT Press (1998)
- Gick, M.L., Holyoak, K.K.: Analogical problem solving. *Cogn. Psychol.* **12**, 306–355 (1980)
- Ho, T.H., Camerer, C., Weigelt, K.: Iterated dominance and iterated best-response in experimental ‘P-beauty contests.’ *Am Econ Rev* **88**(4), 947–969 (1998)
- Kagel, J.H., Levin, D.: The winner’s curse and public information in common value auctions. *Am Econ Rev* **76**, 894–920 (1986)
- Kreps, D.W., Wilson, R.: Sequential equilibria. *Econometrica* **50**, 853–894 (1982)
- Milgrom, P., Roberts, J.: Limit pricing and entry under incomplete information: an equilibrium analysis. *Econometrica* **50**, 443–459 (1982)
- Milgrom, P., Roberts, J.: Adaptive and sophisticated learning in normal form games. *Games Econ Behav* **3**, 82–100 (1991)
- Miller, R.M., Plott, C.R.: Product quality signalling in experimental markets. *Econometrica* **53**, 837–872 (1985)
- Moulton, B.: Random group effects and the prevision of regression estimates. *J Econometrics* **32**, 385–397 (1986)
- Nagel, R.: Unraveling in guessing games: an experimental study. *Am Econ Rev* **85**, 1313–1326 (1995)
- Perkins, D.N., Salomon, G.: Teaching for transfer. *Educ Leadersh.* **46**, 22–32 (1988)
- Robinson, J.: An iterative method of solving a game. *Ann Math* **54**, 296–301 (1951)
- Roth, A.E., Erev, I.: Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. *Games Econ Behav* **8**, 164–212 (1995)
- Salomon, G., Perkins, D.N.: Rocky roads to transfer: rethinking mechanisms of a neglected phenomenon. *Educ Psychol* **24**, 113–142 (1989)
- Selten, R.: “Anticipatory Learning in Games” in *Game Equilibrium Models*, vol. 1, Chap. 3. New York: Springer (1991)
- Stahl, D.O.: Boundedly rational rule learning in a guessing game. *Games Econ Behav* **16**, 303–330 (1996)
- Stahl, D.O.: Action reinforcement learning versus rule learning. *Greek Econ Rev* **22**, 27–56 (2003)
- White, H.: *Estimation, Inference and Specification Analysis*. New York: Cambridge University Press (1994)