

Learning and Transfer in Signaling Games*

David J. Cooper
Department of Economics
Weatherhead School of Management
Case Western Reserve University
10900 Euclid Avenue
Cleveland, OH 44106-7206
djcl3@guinness.som.cwru.edu

John H. Kagel
Department of Economics
Ohio State University
410 Arps Hall
1945 North High Street
Columbus, OH 43210-1172
kagel@economics.sbs.ohio-state.edu

January 26, 2004

Abstract

We explore how learning to play strategically in one game promotes strategic play in related games. Experiment 1 involves a simple linear transformation in payoffs between games (with presentation format changed as well). There is considerable but incomplete transfer as the *growth* in the learning process stalls. Experiment 2 changes responders' payoffs from supporting a pooling equilibrium to supporting a separating equilibrium. More strategic play is observed following the change than for inexperienced subjects in control sessions, contrary to the prediction of a fictitious play model. We present evidence that experience generates increased numbers of sophisticated players who anticipate responders' behavior following the change in payoffs, resulting in positive transfer.

Key words: learning, learning transfer, signaling games, experiment.

JEL classification: 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.

Over the past decade, experimenters have shown that models of learning, in which players have bounded rationality and only gradually learn how to best respond in a game, often capture important features in experimental data that standard equilibrium approaches based on full rationality of players miss. Realizing that learning is a critical element of behavior in games, both theorists and experimenters have produced a rich literature exploring a wide variety of issues about learning and other adaptive processes.¹ One learning issue that has been largely unexplored is the ability of subjects to transfer what they have learned between related games.² Virtually all papers on learning, both theoretical and experimental, employ an environment in which learning takes place within an essentially stationary environment. Yet, even the most casual reflection indicates that many, if not most, real world settings involve a changing environment, so that the ability to take what has been learned in one game and apply it in another, related game is an integral part of learning.

The goal of the present research is to extend the study of learning to encompass transfer of learning between games. Beyond studying whether subjects can transfer what they have learned in one game to a related game, we also explore the mechanism(s) underlying how such transfer takes place. Our work therefore is not only designed to shed light on specific issues of transfer for the games studied, but to illuminate more general issues of how subjects learn across games.

We study these issues in the context of Milgrom and Roberts' (1982) entry limit pricing game which provides a rich environment in which to study learning between games. We report results from two experiments.

Experiment 1 establishes a baseline on subjects' ability to transfer learning between related games by examining what should be an almost trivial transfer from one game to the other: Payoffs in one

¹See Fudenberg and Levine (1998) for a summary of the theory literature on learning. Experimental papers on learning include Camerer and Ho, 1999; Crawford, 1991; Binmore, Gale, and Samuelson, 1995; Brandts and Holt, 1993; Boylan and El-Gamal, 1993; Cheung and Friedman, 1997; Cooper, Garvin, and Kagel, 1997a,b; Cooper and Kagel, 2003; Cooper and Stockman, 2002; Erev and Roth, 1998; Feltovich, 2000; Mookherjee and Sopher, 1997; Roth and Erev, 1995; Roth, Erev, Slonim, and Berman, 2000.

²Previous work on this topic by economists and cognitive psychologists is summarized in Section I of this paper.

game are a linear transformation of payoffs in the other game, with payoffs flipped around the horizontal and vertical axes of the payoff table as well. From a game theoretic point of view the two games are identical, so that if subjects fail to exhibit substantial transfer when crossed-over between these two games, there is little hope for transfer between less closely related games. From a behavioral point of view the games are very similar as well, since strategic play in both involves separating by high quality types (limit pricing by low cost monopolists). We find substantial but incomplete transfer in the cross-over treatment, as the change in games effectively stalls the *growth* in the learning process, compared to the control group, for a number of plays of the game. The net result is that strategic play in the cross-over treatment far exceeds that of inexperienced subjects, but remains behind that of the control group through the end of the experimental session. This result both demonstrates that it is possible for subjects to transfer the ability to play strategically between closely related games and illustrates the difficulty of the task.

Experiment 2 examines a substantially more challenging test of learning transfer between games. In the initial game, receivers' (entrants') payoffs support a pure strategy pooling equilibrium to which inexperienced subjects' play reliably converges. In this equilibrium low quality types (high cost monopolists) act strategically, imitating the high quality types (low cost monopolists). Receivers' payoffs are then changed to eliminate the pooling equilibrium, leaving only pure strategy separating equilibria. While conceptually similar, strategic behavior following this second game requires substantially different actions than in the first game, as it is now the high quality types who must act strategically, distinguishing themselves from low quality types. A fictitious play learning model that has worked well in tracking play from previous signaling games (Cooper *et al*, 1997b) predicts that strategic play by high quality types immediately following the change in responders' payoffs will be *less* than in inexperienced control sessions (negative transfer), and will remain less than the controls until behavior converges to the equilibrium outcome. Contrary to these predictions, high quality types show significantly *more* strategic play immediately following the change in responders' payoffs than in inexperienced control sessions (positive transfer). In fact, the play of subjects following the crossover is statistically indistinguishable

from experienced subjects in the control sessions, suggesting experience with the pooling equilibrium serves as an almost perfect substitute for experience in a game where the only pure strategy sequential equilibria are separating.

Fictitious play's failure to predict the positive transfer observed in Experiment 2 derives from the unsophisticated learning process it incorporates. A fictitious play learner treats his opponents like 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. To capture the rapid jump to strategic play by high quality types observed in the data, we modify the basic model of fictitious play to include sophisticated learners who model how their opponents make decisions, and thereby anticipate the change in responders behavior following the change in their payoffs. Fitting this model to the data, we find a statistically significant fraction of sophisticated learners in the population, *and that the fraction of sophisticated learners increases over time*. With the addition of a growing number of sophisticated learners over time the model tracks the jump in strategic play by high quality types in Experiment 2 after the change in responders' payoffs.

Our experimental results suggest some general conclusions. First, given the results of Experiment 1 there is unlikely to be complete transfer even between closely related games. This suggests the need for learning models that account for the disruptions in learning caused by changing environments. Second, given the results of Experiment 2, it is clear that many subjects are not the simple-minded automata envisioned by most 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. These results also suggest that good models of learning should allow for the *development* of substantial sophistication on the part of subjects over time.³ A third, closely related point is that if the mechanism underlying much of the positive transfer we observe is sophisticated learning, as the data suggests, substantial transfer can still take place even if related games don't require subjects to behave in exactly the same manner.

³For examples of this sort of model, see Milgrom and Roberts, 1991; Selten, 1991; Stahl, 1996, 2000; Camerer, Ho, and Chong, 2002.

Sophisticated learning successfully tracks our data because it allows for subjects who anticipate the behavior of their opponents. To the extent that such strategic anticipation underlies much of game theory, learning to behave in a sophisticated manner should be applicable in many environments beyond the signaling games studied here.

The organization of this paper is as follows: Section I summarizes the previous literature in economics and psychology on learning transfer. Section II describes the limit pricing game and section III describes the experimental procedures. Section IV and V present the results for Experiments 1 and 2 respectively. A model of sophisticated learning is fit to data from Experiment 2 in Section VI. Section VIII summarizes the paper and discusses directions for future research.

I. Previous Work on Transfer

The Psychology Literature: There is a large psychology (and educational psychology) literature dealing with issues of transfer of learning from one situation to another. This literature suggests little reason to predict high levels of positive transfer. In these experiments positive transfer usually only occurs between settings that are perceived as being quite similar. The nature of transfer is often quite narrow, with subjects failing to recognize underlying concepts that allow them to generalize between settings (Gick and Holyoak, 1980; Perkins and Solomon, 1988; Solomon and Perkins, 1989). Further, for transfer to take place subjects must usually either be encouraged to look for similarities between the two situations or be trained in a way that stresses the generality of the concepts being taught (for example, Klahr and Carver, 1988).

Although the psychology literature is not very encouraging, its direct relevance for economic games is questionable. Most psychological 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 or not the factors that promote adjustment *over time* to equilibrium in one game will speed up the adjustment *over time* to a new equilibrium in the same game or in related games. In this framework the fact that games involve strategic interactions may help facilitate transfer, as in adjusting to equilibrium in one game agents are likely to

become sensitized to the strategic behavior of their opponents. This in turn may cultivate either deeper understanding of the underlying strategic contingencies involved in the game and/or sensitize agents to be responsive to the impact of changes in economic contingencies on their rivals' actions. The latter is consistent with the growth in the number of sophisticated learners identified in experiment 2.⁴

The Economics Literature: There have been only limited studies of transfer between games in economic experiments. Kagel and Levin (1986) show that subjects who have learned to avoid the winner's curse in auctions with small numbers of bidders (3 or 4) succumb to the winner's curse when playing with larger numbers of bidders (6 or 7), suggesting that learning is situation specific and does not involve theory absorption. However, they do not employ control groups to determine if prior experience reduced (or perhaps exaggerated) the severity of the winner's curse in auctions with larger numbers of bidders. Kagel (1995) shows that prior experience with first-price sealed-bid common-value auctions reduces the severity of the winner's curse compared to inexperienced bidders in an ascending-price common-value auction. But there is no benefit to prior experience with an ascending-price auction when bidding in a first-price sealed-bid auction. He attributes this difference to the fact that lessons learned in the sealed-bid auctions (bid somewhat conservatively relative to own signal value) are of value for ascending-price auctions, but lessons learned in ascending-price auctions (pay attention to other bidders' drop-out prices) are of no value in sealed-bid auctions.

Ho, Camerer, and Weigelt (1998) study cross-game learning in p-beauty contest games. Subjects are crossed between finite-threshold games ($p > 1$) in which equilibrium can be reached in a finite number of steps of iterated dominance and infinite-threshold games ($p < 1$) in which equilibrium cannot quite be reached in finitely many iterations of dominance. Experienced subjects first round choices are no different than inexperienced subjects. But choices in subsequent rounds converge faster to equilibrium

⁴We are ignoring incentive issues here (economists typically use them, psychologists do not). Incentives generally reduce the variance in outcomes around the mean (see Smith and Walker, 1993, for a review of this literature) and sometimes (but far from always) produce outcomes closer to equilibrium predictions. Our own suspicion is that incentives can induce heightened levels of deductive reasoning in games, which can serve as a partial substitute for experience in games which exhibit systematic adjustment patterns as agents increase their understanding of the problem at hand (Cooper et al., 1999).

than inexperienced subjects, thereby showing some evidence of positive learning transfer.⁵

II) The Limit Pricing Game: The games employed here are based on Milgrom and Roberts' (1982) model of limit entry pricing. This game provides a rich environment for studying transfer. Like most signaling games, it features multiple (sequential) equilibria including pure strategy pooling equilibria and (possibly) pure strategy separating equilibria. This multiplicity allows us to confront subjects with related games that require quite different strategic actions. Further, strategic play is clearly identifiable in the limit pricing game, making it easy to measure the extent to which there is, or is not, cross-game learning.

We begin by laying out a basic version of the game in which players are choosing over quantities. Equilibrium predictions are derived for the quantity game. We finish by describing how the quantity game was transformed into a price game, the basic manipulation underlying Experiment 1.

A. The Quantity Game: Milgrom and Roberts describe a two-stage game with a homogeneous good and a linear market demand curve. The game is played between a monopolist (M) and a potential entrant (E). There are two possible monopolist types, high and low cost (MH and ML). Entry is profitable against an MH but not against an ML. The game begins with M observing its type. The M's cost level is realized according to predetermined probabilities that are common knowledge. In the first period of the game, M chooses a quantity absent any rival producers. E sees this quantity, but not M's type, and either enters or stays out in the second period. It is assumed that if entry occurs the two firms behave as Cournot duopolists in the second period. If entry does not occur, M produces its profit maximizing, uncontested monopoly output in the second period. 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) in the first period.

In our experiment we simplify the game even further, collapsing the two-stage game into a single stage by imposing the second-stage outcomes; the Cournot outcome following entry (IN) or the uncontested monopoly outcome following not entering (OUT). Ms are provided with a payoff table

⁵Also see Larrick, Morgan and Nisbett (1990) who study the transfer effects from teaching normative rules of economic behavior (e.g., ignoring sunk costs).

representing their profit for their first period decision added to the present discounted value of their profit in the second period. E's payoff table reflects the second period returns for staying out or entering and playing a Cournot duopoly game. These changes allow us to focus on the signaling aspects of the game, and have the added benefit of speeding up play, allowing for more repetitions in an experimental session.

Ms' payoffs for the quantity game are given in Table 1a. Ms' choices, 1-7, may be thought of as output levels, with higher outputs corresponding to lower prices.⁶ The prior probabilities of the two M types are 50% throughout and are common knowledge.

[Insert Table 1 here]

Three features of Table 1a capture the main strategic elements confronting Ms. First, all else being equal, Ms are better off if Es choose OUT rather than IN. Second, reflecting their 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 E's behavior -- MLs would choose output 4 as opposed to 2 for MHs. These choices will be referred to as the Ms' "myopic" maxima. Finally, output choices 6 and 7 are dominated strategies for MHs, but not for MLs. At these outputs MLs can, in theory, perfectly distinguish themselves from MHs.

Two different payoff tables, Tables 1b and 1c, were used for Es in the quantity game. These represent "high cost" and "low cost" entrants respectively. Only one of these two tables was in effect at any given time. In both tables it always pays to play IN when M is known to be an MH type and to play OUT against an ML type. However, given the prior probability of the different M types, the expected value of OUT is greater than IN in Table 1b (250 vs. 187) and the expected value of IN is greater than OUT in Table 1c (350 vs. 250).

B. Equilibrium Predictions: For Tables 1a and 1b (the quantity game with high cost Es), there exist multiple pure strategy pooling, as well as separating, equilibria.⁷ Pure strategy pooling equilibria occur at

⁶Payoffs are given in the experimental currency "francs." Francs were converted to dollars with one franc equal to \$.001. Headings in Tables 1 and 2 have been changed to match the exposition in the text.

⁷All of the equilibria to be described are sequential (Kreps and Wilson, 1982).

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 output level 3 rather than deviating to their myopic maxima. Similar out of equilibrium beliefs support the other pooling equilibria.

Pooling equilibria at outputs 3-5 involve 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 6 or always choose output level 7 and never incur entry. With MLs choosing output levels 6 or 7, MHs cannot profitably imitate them as they earn more by choosing their myopic maxima (choice of output 2 dominates outputs 6 and 7 for MHs). Once again the beliefs supporting these equilibria are that any deviation from the proposed equilibrium involves an MH type with sufficiently high probability to induce entry, as 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 Tables 1a and 1c (the quantity game with low cost Es) 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 .80 and 5 with probability .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.*

As is typical of signaling games, the limit pricing game suffers from an overabundance of equilibria. To obtain sharper predictions, we apply the most common of the equilibrium refinements for signaling games, the intuitive criterion of Cho and Kreps (1987). This reduces the equilibria in games with high cost Es to pooling at output 4 or 5, and the efficient separating equilibrium with MLs choosing

6. For games with low cost Es, only the separating equilibrium with MLs choosing 6 survives.⁸

C. Price Game: For Experiment 1 we needed a game that is closely related to the quantity game with low cost entrants. We therefore constructed a comparable game with subjects choosing over “price.” Payoffs in the price game are a linear transformation of the payoffs in the quantity game, with payoffs flipped around the horizontal and vertical axes of the payoff table as well.⁹ These payoffs are shown in Table 2.

[Insert Table 2 here]

To get Ms’ payoffs in the price game, subtract 50 from the payoffs in Table 1a and then multiply by 1.25. The payoff tables are then flipped from top to bottom, and the location of low cost types’ payoffs are flipped from right to left. Es’ payoffs are obtained from Table 1c by subtracting 25 from low cost entrants payoffs and then multiplying by 1.25. The position of the two columns are then flipped. From a game theoretic point of view, the price game and the quantity game are identical. None of the equilibrium predictions are affected by the transformation (once we control for the flipping of the payoff tables), nor are the incentives necessary to induce strategic play affected.

III) Experimental Procedures and Design: We begin this section by describing the general procedures used in all experiments. We then lay out the specifics of the experimental design. Table 3 summarizes the main details of the experimental treatments.

[Insert Table 3 here]

General Procedures: Each experimental session 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

⁸For example, consider the inefficient separating equilibrium with MLs choosing 7. Roughly, the intuitive criterion tells us that out of equilibrium beliefs should not put positive weight on deviations that can’t conceivably increase a player’s payoffs over his payoffs in the candidate equilibrium. Suppose a 6 is played. This deviation can never be profitable for a MH since 6 is a strictly dominated strategy for them. Therefore the beliefs must put 100% weight on this deviation coming from an ML. It follows directly that playing 6 is a profitable deviation for an ML, destabilizing the separating equilibrium where MLs choose 7.

⁹Flipping alone would be totally transparent. A linear transformation without flipping, while somewhat less transparent, doesn’t force players to change actions in order to continue playing strategically. Without forcing changes in actions it is difficult to determine whether any meaningful transfer occurs, as opposed to simple inertia. Hence, we made both changes.

were read out loud, with each subject having a written copy.¹⁰ By reading the instructions out loud, we ensure that details of the game are common knowledge. For example, all subjects hear (and know that the other subjects hear) that the ex ante probability of each M type is 50%. 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.

Before each play of the game the computer randomly determined each M's type and displayed this information on the Ms' screens. Ms chose first, with each M's choice sent to the E they were paired with for that game. The program automatically highlighted the possible payoffs for each choice made by the Ms and asked them to confirm their choice, limiting the possibility of mistakes. Es then decided between IN and OUT. Once again, possible payoffs were highlighted and subjects were asked to confirm their decisions. 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 portion of each subject's screen displayed the results of all pairings. In particular, we displayed for each pairing M's type, M's action, and E's response. Thus, subjects had access to a full history of Ms' actions conditioned on their type and Es' responses conditioned on Ms' actions. Subject identification 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.¹¹ Experienced subject sessions had a minimum of 36 games, with all but two of the control

¹⁰A copy of the instructions is available on Cooper's website, www.weatherhead.cwru.edu/djcooper.

¹¹One session had only 24 games since it was conducted in an undergraduate economics class during class time, which limited the number of games.

sessions having 48 games. All of the crossover treatments use experienced subjects.

When crossovers 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 (e. g., Ms continued to choose first, followed by Es responses) but that payoffs had changed. The number of additional games to be played was announced as well.

Subjects were recruited through announcements in undergraduate classes and 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.¹²

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. Econometric analysis indicates that there are no systematic differences between choices in the inexperienced sessions for subjects who later returned for an experienced subject session and those who did not.

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

¹²These sessions also tended to be shorter since only an abbreviated version of the instructions were read and subjects were familiar with the game.

switched between generic context and meaningful context or vice versa. For sessions with a crossover from the quantity to price game (or vice versa), labels were changed following the crossover where appropriate. For example, in a meaningful context session the strategies would change from being labeled as "output" to being labeled as "price." Any such changes were noted in the instructions subjects received for the cross-over.

Earlier results find that the use of meaningful context speeds up learning for inexperienced subjects (Cooper and Kagel, 2003a), but does not affect the play of experienced subjects. We employ control variables for any potential context effects in the statistical analysis.¹³

Experimental Design, Experiment 1: Experiment 1 crosses subjects from the quantity game with low cost entrants (Tables 1a and 1c) to the price game (Tables 2a and 2b) or vice versa. The relatively easy crossover in Experiment 1 should provide a best-case scenario for generating positive transfer. The flip side of this is that if subjects fail to exhibit substantial transfer between these two games there would be little point to exploring transfer between less closely related games.

The data set for Experiment 1 includes 131 subjects who returned for an experienced subject session; 65 subjects in five crossover sessions and 66 subjects in five control sessions without a crossover.¹⁴ Experienced subjects in the crossover sessions were switched from the quantity game with low cost entrants to the price game (or vice versa) following 12 periods of play. Play then continued for an additional 36 periods.

Experimental Design, Experiment 2: Experiment 2 pushes subjects ability to learn between games by confronting them with a more difficult crossover. Subjects in Experiment 2 are first trained in the limit

¹³Due to the limitations of the data set, we cannot identify any interaction between the type of context employed and the size of the crossover effect. This can be seen most clearly for Experiment 2, where all of the crossover sessions employed a meaningful context. 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. Given the lack of any crossover sessions using generic context, we cannot determine if varying the context would affect the size of the crossover effect. This issue is a topic for future research.

¹⁴One crossover session included thirteen subjects who had already been in an earlier experienced session. Data generated by these twice-experienced subjects is not included in our analysis, nor are these subjects included in the count of 65 subjects in the crossover sessions. Completely excluding the data from this session has no effect on our qualitative conclusions.

pricing game with high cost Es (Tables 1a and 1b). Past experience with this game shows that play will reliably converge to the pooling equilibrium at output 4 (Cooper, Garvin, and Kagel, 1997b). The crossover treatment in this game involves switching Es payoffs from Table 1b to Table 1c. This change eliminates all pooling equilibria. We compare the development of strategic play by MLs (play of output levels 5 - 7) following the crossover to the control group where Es are low cost types throughout.

The design of Experiment 2 is driven by past analysis of games of this sort. Models based on fictitious play have done a good job of tracking the evolution of play in earlier experiments with this limit-pricing game (Cooper, Garvin and Kagel, 1997a,b; Cooper et al., 1999). In Section VI we develop the predictions of such a model for Experiments 1 and 2 at length, but for now a summary will suffice. For the limit pricing game with high cost Es, the fictitious play model (without sophisticated learners) predicts that inexperienced subjects will pool at output 4, with minimal strategic play by MLs. The model also predicts that, following a crossover from high cost Es to low cost Es, the probability of strategic play by MLs will be less than or equal to the probability of strategic play as MLs for inexperienced subjects. In other words, the fictitious play model predicts (weakly) negative transfer. Intuitively, fictitious play learners treat their opponents as a fixed distribution. Unable to anticipate the effect of changing their opponents' payoffs, these unsophisticated learners have no reason to switch to strategic play as MLs following the crossover. Only when MLs have had an opportunity to learn from direct observation that Es' behavior has changed following the crossover will they begin to limit price.

One prediction of the fictitious play model (without sophisticated learners) is that the degree to which MLs play strategically following the crossover in Experiment 2 will be adversely affected by the extent of their experience with the high cost entrant game. To determine if any such effect exists, 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 second (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.¹⁵ In the third session half the subjects had played in one prior session with high cost Es and half had played in two prior sessions with high cost Es. This session was crossed to the low cost E treatment in the 25th game.

The control (no crossover) data set for experiment 2 is the same as the control in Experiment 1.

IV) Results for Experiment 1: Crossovers Between the Quantity and Price Games: To ease the exposition, all data have been transposed to the quantity game regardless of whether it comes from the quantity or the price game. For example, if an M in the price game chooses price level 2, this gets transposed into output level 6. Empirically, behavior is indistinguishable between the two payoff tables prior to the crossover.¹⁶

[Insert Figure 1 here]

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. Strategic play by MLs increases steadily over time. By the end of the experienced control sessions, the vast majority of MLs are playing strategically.

To give a frame of reference for the crossover sessions, we start by seeing how play develops in the control sessions. Figure 1 reports the relative frequency of Ms' choices at each output level along with the entry rates in 12 period cycles. For inexperienced sessions, this figure only includes data from subjects who later returned for an experienced session.

In the first twelve periods play by both M types is clustered around their respective myopic maxima (output 2 for MHs, 4 for MLs). Responding to the high entry rates at low output levels, MHs attempt to pool with MLs by choosing 4. By the final 12 periods of inexperienced play, more than half of MHs' choices involve strategic play (choice of 3 and 4). Given that the entry rate in these final 12 periods is much lower for 3 and 4 than for 1 and 2, this movement by MHs is well justified.¹⁷ As MHs

¹⁵One subject was once experienced in this session. She was needed to make an even number of players.

¹⁶Formal statistical tests of this statement are described in footnote 35 in the appendix.

¹⁷An entry rate differential of only 13% is required to make 4 more profitable for MHs than 2 or to make 6 more profitable for MLs than 4.

begin to imitate MLs, the entry rate for output 4 rises. Given the lower entry rates on outputs 5 and 6, strategic play at either of these outputs is more profitable for MLs than choosing the myopic maximum in the final two cycles of the inexperienced control sessions. None the less, the myopic maximum remains the modal choice for MLs in the final twelve period cycle of inexperienced subject play.

The first twelve period cycle of the experienced control sessions looks much like a continuation of the inexperienced sessions. Entry rate differentials continue to support strategic play by both MHs and MLs, strategic play by MLs continues to increase, but the myopic maximum (output level 4) remains the modal choice for MLs. Over time, incentives to limit price become stronger and stronger for MLs, and strategic play by MLs steadily increases. By the second twelve period cycle of the experienced control sessions, play of the myopic maximum is no longer the mode for MLs, and by the final two cycles (periods 25 - 36 and 37 - 48) the myopic maximum is getting less play than either output level 5 or 6. To summarize, the dynamic in the control sessions shows a slow but steady movement by MLs towards strategic play.

In considering how this dynamic is disrupted by the crossover in Experiment 1, we focus on strategic play by MLs. We focus on MLs rather than MHs for several reasons: (1) unlike MHs, the use of outputs greater than the myopic maximum by MLs unambiguously involves strategic play, (2) this strategic play is consistent with an equilibrium for MLs (whereas choices of 3 and higher are not for MHs), and (3) the slow emergence of strategic play by MLs in the control sessions indicates that it is more difficult to master than MHs' early efforts to mimic MLs.

Conclusion 2: The crossover from the quantity game to the price game (and vice versa) retards the development of strategic play for MLs compared to the control group. Strategic play does not revert back to inexperienced subject levels, rather it remains flat at the pre-crossover level for a full cycle of play, and then begins to grow again.

[Figure 2 here]

To illustrate the effect of the crossover on MLs, Figure 2 graphs the percentage of strategic play by MLs in each treatment over time. Choices of output levels 5, 6, and 7 are classified as strategic play by MLs since these choices involve quantities greater than the myopic maximum. Prior to the crossover,

the two graphs look like twins. In both cases, there is a steady increase in the proportion of strategic play. The small gap between the treatments is roughly constant over time. At the point of the crossover, a large gap emerges between the two treatments. While strategic play by MLs continues to grow steadily in the sessions without a crossover, strategic play by MLs plateaus in the sessions with a crossover. Unlike the control sessions, the myopic maximum remains easily the modal choice for MLs immediately following the crossover (periods 13-24 of the experienced crossover sessions). Although growth eventually resumes in the crossover sessions, the level of strategic play by MLs never catches up with the level in the control sessions.¹⁸

This crossover effect can be more closely examined by breaking the data down at the finest possible level, looking at each time a subject plays as an ML.¹⁹ There is steady growth in strategic play immediately prior to the cross-over with the proportion of strategic play *increasing* by an average of 4.6% per play for the three plays (as MLs) immediately prior to the cross-over. In contrast, for the first play as an ML immediately following the crossover the percentage of strategic play *decreases* by 3.1%. It then recovers to its former growth rate in subsequent periods, increasing by an average of 5.6% per play in the following three plays as an ML. In other words, the stall in learning is coming from a quick drop in strategic play by MLs followed by a gradual recovery.

Although the results for Experiment 1 have been presented in terms of descriptive statistics, formal regression analysis demonstrates the statistical significance of the pause in MLs' learning following the crossover. The regression analysis controls for individual effects in the data (repeated observations for the same individuals) and the behavior of Es. The results of this analysis show that the

¹⁸The crossover has virtually no effect on MHs' choices as compared to the control sessions, but does cause a small (but statistically significant) change in Es behavior. This change in Es behavior reduces the incentives for MLs to limit price, a factor that is controlled for in the formal econometrics of the appendix.

¹⁹ This is as detailed a play by play account of strategic behavior by MLs as is possible under our design since subjects switch roles between E and M every six games, as well as changing their type of M between games. Thus, if we were to take the three periods prior to the crossover we would have an incomplete sample of subjects choices as half the time they are playing as MHs (or as Es). By selecting out for each subject each of the three periods prior to the crossover they played as MLs we have a complete sample of their play as MLs.

stall in learning by MLs following the crossover cannot be attributed to changes in entry rates following the crossover.²⁰ The statistical analysis also indicates that this stall in strategic play is not compensated for over the last two cycles of play, so that strategic play by MLs in the cross-over group is below that of the controls through the end of the experienced subject sessions. Although the formal statistical analysis is essential to the conclusions reached, it is technical in nature and does not change any of the conclusions we reach from examining the descriptive statistics. As such it is reported in an appendix.

V) Results from Experiment 2: Many plausible stories can be told about the mechanism underlying positive transfer in Experiment 1. At one extreme, a subject may be learning to limit price through a trial and error learning process without any deeper understanding of the nature of the game. This subject can nevertheless transfer what they have learned in the quantity game and successfully apply it to the price game by recognizing the relatively straight forward mapping between the two games. Once the mapping is recognized, strategies that worked well before will be successful following the crossover. Thus, positive transfer in Experiment 1 only requires minimal understanding of the strategic elements of the game.

At the other extreme a subject may have mastered the general principles underlying the game, allowing them to easily master the crossover from the quantity to the price game. This sort of ultra-sophisticated learner not only learns to play strategically, but also forms a model of how their opponents are playing (and why) and understands why playing strategically is a good response to Es' behavior. For successful transfer between the price and quantity games, this subject needs to recognize that the same general principles apply to both games, which should be relatively easy given the strong similarity between the two games.

One can easily imagine many intermediate levels of sophistication between these two extremes. The point is that these two strongly contrasting approaches, learning an action and learning a general principle, yield observationally equivalent predictions for Experiment 1. Experiment 2 examines a

²⁰The statistical analysis also controls for the fact that several sessions employed meaningful context, whereas others did not.

crossover that allows us to distinguish between differing levels of sophistication in subjects' learning.

Thus, it enables us to better understand the behavioral mechanism underlying transfer between games.

Conclusion 3: In games with high cost Es, play converges towards a pooling equilibrium at output level 4, with little strategic play by MLs at any point in time.

[Insert Figure 3 here]

Figure 3 shows inexperienced subjects' play in the game with high cost Es, and the last cycle of experienced subjects' play prior to the change in Es' payoffs.²¹ For the first twelve period cycle, the myopic maximum is the modal outcome for both M types. However, MLs and MHs face quite different incentives to move from their myopic maxima to strategic play. For MHs, the entry rate differential between outputs 2 and 4 is 36.5% for the first twelve period cycle, well above the 13% differential needed to make strategic play incentive compatible, and grows to 59.3% for the final twelve period cycle. For MLs, the entry rate differential between 4 and 6 hovers at 19% in the first twelve period cycle, just a little above the 13% differential needed to make strategic play incentive compatible, and drops to 12% by the final twelve period cycle. Consistent with their strong incentives to play strategically, a large proportion of MHs choose output 3, 4 or 5 even in the first twelve period cycle (47.1%), with the frequency of strategic play growing steadily over time after that. For the final twelve period cycle of the inexperienced play, output level 4 is easily the modal outcome for MHs (53.5%), and 63.2% of MH play can be classified as strategic play (outputs 3, 4 or 5). Consistent with their weak incentives to limit price, there is little strategic play by MLs even in the first twelve period cycle (9.2%), with the proportion of strategic play changing only slightly over time (14.9% in the final 12 period cycle). Overall, there is smooth movement towards the pooling equilibrium at 4.

In the last twelve period cycle before Es payoffs changed, differences in the incentives to limit price have become even more extreme.²² The entry rate differential between 2 and 4 has risen to 72.4%,

²¹Only subjects who participated in the cross-over sessions are included in Figure 3.

²²Because subjects are crossed over at different times, this panel pools subjects with differing levels of experience. The purpose here is to provide a snapshot of subjects' behavior immediately prior to the crossover.

while the entry rate on 4 has fallen to 7.6%.²³ Following these strong incentives, the pooling equilibrium at 4 has gained even more strength. For MHs, 60.2% of play is at 4, and 67.3% involves strategic play of some sort. For MLs, 89.6% of play is at 4 and there is almost no strategic play (5.2%).

Conclusion 4: MLs in the crossover treatment look much 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.

In considering the effect of the crossover from high to low cost Es our focus is once again on strategic play by MLs. Following the crossover, it immediately becomes incentive compatible for MLs to play strategically – the entry rate differential between output 4 and 6 jumps to 44.8% in the first cycle following the crossover. Strategic play by MLs in the first cycle following the crossover is substantially higher than in the first cycle of inexperienced subject play in the control group (25.7% versus 8.5%). It is also substantially higher than for these same MLs in the last cycle prior to the crossover (25.7% versus 5.2%). MLs following the crossover are neither continuing to do what they did before the crossover, nor are they behaving like naive subjects. Indeed, inexperienced subjects in the control group never achieve the level of strategic play observed for MLs in the first twelve period cycle following the crossover. It is only by the first twelve period cycle of the *experienced* subject control group that we see more strategic play by MLs (40%) than immediately following the crossover. Thus, experience in the game with high cost Es appears to be a good substitute for experience in the game with low cost Es.

[Figure 4 here]

Figure 4 provides a 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 the control sessions, time is measured from the beginning of the session. For example, “Play 1” is the first time an inexperienced subject has played as an ML. In the crossover sessions time is measured from the point of the crossover. In this case “Play 1” is the first time a subject has played as an ML following the crossover. The graph plots the percentage of strategic play by MLs in

²³Given that there are no choices of output 6 in this cycle, we cannot actually calculate the entry rate differential between outputs 4 and 6. However, it can't be greater than the 7.6% entry rate for output 4.

inexperienced control sessions, experienced control sessions, and in the crossover sessions following the crossover to low cost Es. Looking at “Play 1,” MLs in the crossover treatment immediately limit price more often than their counterparts in the inexperienced control sessions. This suggests that MLs are anticipating 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. To summarize, subjects following the crossover look much more like experienced than inexperienced subjects in the control sessions. Experience in the game with high cost Es helps subjects learn to play strategically as MLs in the game with low cost Es.

One potential confounding factor here is that there was somewhat more incentive for MLs to limit price following the crossover than in the control sessions.²⁴ 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 formal statistical analysis also 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 further evidence against fictitious play learning, absent the presence of a sizable number of 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 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 may also reflect deeper aspects of subjects’ learning. The structural model of learning developed in Section VI 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

²⁴For the first twelve period cycle following the crossover, the entry rates following output levels 2, 4, and 6 are .978, .615, and .167 respectively. The corresponding figures are .800, .418, and .333 for the first cycle of the inexperienced control sessions and .917, .609, and .182 for the first cycle of the experienced control sessions.

likely to play strategically for the game with high cost entrants and the game with low cost entrants, thereby helping to generate the positive correlation between individual subjects' strategic play prior to and after the crossover.

Finally, it is worth contrasting the results here with those of Experiment 1. In Experiment 1, the cross-over treatment did not require any change in strategic behavior following the cross-over - subjects should limit price as MLs both before and after the crossover. The cross-over disrupted MLs' strategic play relative to before the cross-over and relative to the control group - essentially stalling the growth in strategic play for the cross-over group for a full cycle of play. In contrast, in Experiment 2 the cross-over takes subjects who had learned to play strategically as MHs (with no strategic play in their role as MLs) and requires that they change their behavior, playing strategically as MLs. This would seem substantially more demanding than the task confronting subjects in Experiment 1. However, in contrast to what one might naively project on the basis of the stall in strategic play of MLs following the crossover in Experiment 1, in Experiment 2 strategic play of MLs was substantially higher than either before the cross-over and substantially higher than for the inexperienced control group.

VI. A Structural Model of Learning and Sophistication: The data from Experiment 2 indicate that subjects play far more strategically as MLs following the crossover than inexperienced subjects do in the control sessions. This result is inconsistent with the predictions of the fictitious play model that motivated Experiment 2. Since these predictions were driven by the inability of Ms to anticipate changes in Es' behavior following the crossover, their failure suggests that some Ms are *at least* sophisticated enough to anticipate the change in Es' behavior.

To explore this intuition more formally, this section develops and fits a structural model of stochastic fictitious learning, examining the implications of adding sophisticated learners to the model. The purpose of this analysis is threefold. First, we want to show that the addition of sophisticated learners improves the econometric fit to the data from Experiment 2. 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 in Experiment 2. Finally, we want to confirm that the

learning model *without* sophisticated learners is sufficient to explain the behavior of MLs following the crossover in Experiment 1. While there are presumably sophisticated learners present in the data from Experiment 1, it is only when we consider the more challenging crossover in Experiment 2 that their presence becomes necessary to track the data.²⁵

A. The Learning Model: 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 similar models have done a good job of tracking the evolution of play in earlier signaling game experiments (Cooper, Garvin, and Kagel, 1997a, b; Cooper et al., 1999).²⁶ The model, although only described for Ms in our limit pricing game, generalizes in a straight forward way to other games.

A belief-based learning model requires an algorithm for choosing a strategy in period t given beliefs, a rule for updating beliefs from period t to period $t+1$, and a rule for 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 roughly 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 formulas:

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

²⁵The econometric analysis in the appendix employs a non-structural model of behavior. The work there is designed to provide statistical support for the patterns identified in the text while controlling for the obvious covariates that could, potentially, confound the results reported, and is not designed to test any specific model. In contrast, this section asks if the surprising degree of cross-game learning observed in Experiment 2 can be explained by the presence of, and increasing numbers of, sophisticated learners. The answer to this question requires the development of a structural learning model.

²⁶We have not explicitly considered other classes of learning models such as the reinforcement learning model of Roth and Erev (1995) or the EWA model of Camerer and Ho (1999). Determining the learning model that best tracks the behavior of subjects goes well beyond the scope of the present paper. Our goal is to develop a learning algorithm that can track behavior in Experiment 2. Given that these other learning models embody unsophisticated learners, they would presumably require some sort of similar augmentation to track the crossover in Experiment 2.

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

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.²⁷ Player i 's probability of choosing output j in period t (subject to choosing a new output), p_{ij}^t , is then generated via a logit rule:²⁸

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

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 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 equation 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 isn't playing as an M .

²⁷In previous fitting exercises (Stahl, 2000; Cooper and Stockman, 2000), introducing autocorrelation into the model substantially improved the fit.

²⁸The probability of a player choosing the same output twice in a row (as the same type) is greater than $1 - p_{\text{change}}$ since he may reselect the same output when picking a new strategy.

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

For periods following a crossover, the updating rule needs to account 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 $C'_{ij}{}^t(\text{OUT})$.

$$C'_{ij}{}^t(\text{IN}) = (1 - \rho)C_{ij}^t(\text{IN}) + \rho C_{ij}^0(\text{IN}) \quad (\text{eq. 4})$$

Intuitively, $(1 - \rho)$ gives the weight subjects put on experience from the previous related game. Setting $\rho = 1$ is equivalent to saying that there is no cross-game learning and setting $\rho = 0$ is equivalent to saying that the games are treated as being identical.²⁵

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 7 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 types: non-learners and sophisticated learners. Non-learners start with the same initial beliefs as unsophisticated learners, and 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 are maximizing their payoffs subject to their beliefs. This implies that a sophisticated learner anticipates that changes in

²⁵We have explored a variety of other specifications for how beliefs might be transformed following the crossover. The qualitative results are unaffected by alternative specifications.

payoffs will affect Es' choices, and that Es' behavior will change as they accumulate more 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 reported in Experiment 2. The type 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, Ho, and Chong (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 Brants 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 unsophisticated Es' beliefs. Thus, 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.²⁶

The model allows players to switch types 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.²⁷

B. Fitting the Model: The model was fit using data from all subjects (including the controls) who returned for an experienced subject session in Experiment 2. 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 these parameters, 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 in games with low cost Es equal, as a log likelihood ratio test fails to reject the null hypothesis of identical initial beliefs ($\chi^2 = 9.68$, 7 d.f., $p > .10$). 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 the following output classes: low outputs (1 and 2), intermediate outputs (3 and 4), and high outputs (5, 6, and 7). Relaxing these restrictions would only strengthen our main conclusions.

²⁶The 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 process.

²⁷We do not assign specific types to the subjects. Instead, we generate the likelihood of a subject’s observed choices subject to being a certain type, and then generate 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.

The data set includes repeated observations from the same individuals which cannot be treated as being 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 apply the correction for clustering suggested by Moulton (1986) and White (1994) to the standard errors. This correction is a close variation on the commonly used sandwich estimator of the variance-covariance matrix. As it turns out, this correction has only a small effect on the standard errors for the parameter estimates of primary interest.

Conclusion 5: The addition of “sophisticated” learners to the basic model of fictitious play generates a statistically significant improvement in the fit to the data.

[Table 4 here]

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 for all models 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 an enormous improvement in the log-likelihood ($\chi^2 = 603.60$, 5 d.f., $p < .01$). Not surprisingly, given the strong dynamics in the data, the evidence in favor of learning is overwhelming. The improvement in the log-likelihood going from Model 2 to Model 3 is also enormous and significant at the 1% level ($\chi^2 = 189.42$, 6 d.f., $p < .01$). 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 statistically significant, the population is clearly moving toward greater sophistication over time.²⁸

C. Simulations: The simulations reported in this section show that the learning model without sophisticated learners misses important features of the data for Experiment 2 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.

Conclusion 6: The learning model with sophisticated learners generates substantially better tracking of MLs' behavior in the crossover treatment of Experiment 2. Growth over time in the proportion of sophisticated learners provides a mechanism for the positive transfer observed in the crossover treatment of Experiment 2.

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 of Es and MHs are generated randomly using the observed frequencies in the data.²⁹ 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), using the fitted probabilities of each type to generate aggregate behavior.

[Figure 5 here]

²⁸Some secondary features of the parameter estimates in Model 3 are worth noting. First, the estimate for ρ , the "reset" parameter, is not significantly different from 0 at even the 10% level. Unsophisticated learners' beliefs are barely affected by the crossover. Second, the estimate of "weight on others' experience" is tiny and statistically insignificant. The beliefs of unsophisticated learners are based almost entirely on their own experience. Finally, the probability of changing strategies is always significantly less than 1, indicating substantial autocorrelation.

²⁹For each twelve period cycle for each treatment (no crossover vs. crossover), we calculated the probability of each output level being played and the probability of entry following each output level. These probabilities were used to randomly generate the choices of Es and MHs observed by the simulated MLs.

Figure 5 displays strategic play by MLs from the simulations in the same way that Figure 4 did for the experimental data. The unit of time on the x-axis is how many times a subject has played as an ML. As in Figure 4, each panel plots the percentage of strategic play by MLs in inexperienced and experienced control sessions, and in the crossover sessions. The top panel reproduces the data from the experiment (Figure 4), the middle panel simulates play without sophisticated learners, and the bottom panel simulates play with sophisticated learners.

Comparing the top and middle panels of Figure 5, 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 fit the rapid jump to strategic play following the crossover is by allowing for very fast learning following the crossover. (This is the reason for the higher values of the discount parameter δ and the reset parameter ρ in the model without sophisticated learners, as compared to the model 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 from Experiment 2.

The graph in the bottom panel, showing data from the simulations with sophisticated learners, looks much more like the experimental data: Simulated subjects in the crossover treatment immediately show more strategic play as MLs following the crossover than simulated inexperienced subjects in the control treatment. Following the crossover, strategic play by simulated MLs grows gradually, paralleling the growth of strategic play for experienced simulated 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 explains why subjects should play strategically as MLs at least as much following the crossovers as in the inexperienced control sessions. However, the presence of sophisticated learners isn't sufficient to explain why there is *more* strategic play by MLs following the crossovers than in the inexperienced

control sessions. Following the crossover in Experiment 2, sophisticated learners immediately jump to strategic play as MLs. However, they are only somewhat more likely to play strategically (due to greater experience and greater incentives for strategic play) than sophisticated learners in the inexperienced control sessions. If there weren't any more sophisticated learners in the population following the crossover than in inexperienced control session, we would see almost exactly the same level of strategic play as in the inexperienced control sessions.³⁰ 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 related game*. That is, the big difference between the crossover sessions and the inexperienced control sessions is that previous experience with the high cost entrant game results in a higher percentage of sophisticated learners in the subject population. Thus, the primary mechanism underlying the surprising degree of positive transfer in Experiment 2 is the growth in sophistication in the subject population.

[Figure 6 here]

We claimed earlier that the basic model of fictitious play, not including sophisticated learners, was sufficient to predict MLs behavior in Experiment 1. To substantiate this claim we fit the learning model with non-learners and unsophisticated learners to the control data from Experiment 1 and simulate MLs' behavior in the crossovers using these parameter values.³¹ We ran sets of 10,000 simulations for different values of ρ , the reset parameter. Figure 6 shows strategic play by MLs in these simulations as ρ varies. Again, the unit of time on the x-axis is how many times a subject has played as an ML. In the control sessions this is measured from the beginning of the session, while in the crossover sessions it is measured from the point of the crossover with 1 denoting the first play following the crossover. By

³⁰Specifically, suppose we reran the simulations of Model 3 but didn't allow 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 in Experiment 2, 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.

³¹The fitted values of the salient parameters with standard errors are as follows: $\lambda = .019$ (.001), $p_{\text{change}} = .715$ (.017), $\delta = .086$ (.010), $w_{\text{other}} = 2.4 \times 10^{-2}$ (4.3×10^{-2}), and probability, non-learner = .230 (.060).

varying the value of ρ , we can easily generate the sort of slowdown seen in the data. For example, with $\rho = .05$ (a very modest reset to inexperienced subject play) the level of strategic play by MLs following the first twelve period cycle after the crossover is almost exactly the same as immediately prior to the crossover. Sophisticated learners are presumably present in Experiment 1 (indeed, a statistically significant proportion of sophisticated learners can be detected using the control data alone), but their presence is not necessary to organize the data. This is in sharp contrast to Experiment 2 where their presence plays a critical role in organizing the data.

Conclusion 7: While sophisticated learners are no doubt present in Experiment 1, their presence is not needed to explain the crossover effects observed in Experiment 1.

VII) 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, but not all, of this knowledge to related games even if the actions necessary to play strategically are quite different. More important than establishing the presence of positive transfer, 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 to the strategic implications of the game that allows them to perform so well compared to naive subjects when put into a new (but related) setting.

Although we attribute the positive transfer in Experiment 2 to the existence of a growing population of sophisticated learners, the experimental design described above 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. All we can state with certainty is that a learning model with sophisticated types provides a consistent explanation for the observed data while one without sophisticated types cannot, and that we have not found a plausible alternative explanation. However, in subsequent research in which two person teams play the roles of monopolist and potential entrants, 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, 2003b).

It is worth noting that our results are far more encouraging regarding the possibility of positive transfer than those typically reported by psychologists. This no doubt reflects differences in the tasks subjects are confronted with and the environments in which subjects are allowed to learn. The psychology literature has tended to focus on learning specific skills (e.g. how to drive a truck) or how to solve certain classes of problems (e.g. logic puzzles). Games, by their very nature, are interactive. As a result, a large part of what is learned is an understanding of how others will behave and, as shown here, this understanding can aid in transferring learning from one environment to another. This interactive element of games is missing from the individual choice problems typically studied by psychologists.

We have only begun to scratch the surface of cross-game learning issues as they apply to economic environments. All of the crossovers used in this paper take place within a relatively narrow class of games. Finding positive transfer here, especially in the more difficult circumstances of Experiment 2, would seem to be necessary, but far from sufficient, to inspire confidence in similar levels of transfer for real economic environments. On the plus side, the interactive thinking that plays a central role in generating the positive cross-game learning identified here is present in most strategic settings. On the negative side, the experimental design we use makes it relatively easy to see that there is a relation between the different games being played. In field settings, agents both have to recognize the underlying game structure and then make the connection between related structures. Because recognizing this connection may be a significant stumbling block, exploring transfer will require directly addressing

questions such as context effects, framing effects, and playing in teams rather than as individuals, topics economists have traditionally been reluctant to explore. The true level of sophistication reached by subjects is also an open question. Experiments where the games are less closely related, or unrelated games are played between the games of direct interest, may also help shed light on this issue. These items are on our agenda, but unfortunately their exploration lies well beyond the scope of the present inquiry.

References

- Binmore, Ken, John Gale, and Larry Samuelson (1995), "Learning to be Imperfect: The Ultimatum Game," *Games and Economic Behavior*, 8, 56-90.
- Boylan, Richard and M. El-Gamal (1993), "Fictitious Play: A Statistical Study of Multiple Economic Experiments," *Games and Economic Behavior*, 5, 205-222.
- Brandts, Jodi and Charles A. Holt (1992), "An Experimental Test of Equilibrium Dominance in Signaling Games," *American Economic Review*, 82, 1350-1365.
- Brandts, Jordi and Charles A. Holt (1993), "Adjustment Patterns and Equilibrium Selection in Experimental Signalling Games," *International Journal of Game Theory*, 22, 279-302.
- Camerer, Colin F., and Tech Ho (1999), "Experience-Weighted Attraction in Games," *Econometrica*, 67, 827-874.
- Camerer, Colin F., Tech H. Ho and Juin-Kuan Chong (2002), "Sophisticated Experience- Weighted Attraction Learning and Strategic Teaching in Repeated Games," *Journal of Economic Theory*, 104, 137 – 188.
- Cheung, Yin-Wong and Dan Friedman (1997), "Individual Learning in Normal Form Games: Some Laboratory Results," *Games and Economic Behavior*, 19, 46-76.
- Cho, In-Koo and David M. Kreps (1987), "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102, 179-221.
- Cooper, David J., Susan Garvin and John Kagel (1997a), "Adaptive Learning vs. Equilibrium Refinements in an Entry Limit Pricing Game," *Economic Journal*, 107, 553-575.
- Cooper, David J., Susan Garvin and John Kagel (1997b), "Signaling and Adaptive Learning in an Entry Limit Pricing Game," *Rand Journal of Economics*, 28, 662-683.
- Cooper, David J. and John H. Kagel (2003) (a), "The Impact of Meaningful Context on Strategic Play in Signalling Games," *Journal of Economic Behavior and Organization*, 50, 311-37.
- Cooper, David J. and John H. Kagel (2003) (b), "Are Two Head Better than One? Team versus Individual Play in Signaling Games." (in preperation).
- Cooper, David J., John H. Kagel, W. Lo and Q.L. Gu (1999), "Gaming Against Managers in Incentive Systems: Experimental Results with Chinese Students and Chinese Managers," *American Economic Review*, 89, 781-804.
- Cooper, David J. and Carol Kraker Stockman (2002), "Fairness and Learning: An Experimental Examination," *Games and Economic Behavior*, 41, 26 – 45.
- Crawford, Vincent P. (1991), "An Evolutionary Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games and Economic Behavior*, 3, 25-59.

- Erev, Ido and Alvin E. Roth (1998), "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique Mixed-Strategy Equilibria," *American Economic Review*, 88, 848-881.
- Feltovich, Nick (2000), "Reinforcement-Based vs. Beliefs-Based Learning Models in Experimental Asymmetric-Information Games," *Econometrica*, 68, 605-642.
- Fudenberg, Drew and David M. Kreps (1988), "A Theory of Learning, Experimentation, and Equilibrium in Games," Mimeo, MIT.
- Fudenberg, Drew and D. K. Levine (1995), "Consistency and Cautious Fictitious Play," *Journal of Economic Dynamics and Control*, 19: 1065-1090.
- Fudenberg, Drew and David Levine (1998), *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Gick, Mary L. and Keith K. Holyoak (1980), "Analogical Problem Solving," *Cognitive Psychology*, 12, 306-355.
- Ho, Teck-Hua, C. Camerer, and K. Weigelt (1998), "Iterated Dominance and iterated Best-Response in Experimental 'P-Beauty Contests,'" *American Economic Review*, 88:4, 947-969.
- Kagel, John H. (1995), "Cross-Game Learning: Experimental Evidence from First-Price and English Common Value Auctions," *Economics Letters*, 49, 163-170.
- Kagel, John H. and Dan Levin (1986), "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, 76, 894-920.
- Klahr, David and Sharon W. Carver (1988), "Cognitive Objectives in a LOGO Debugging Curriculum: Instruction, Learning, and Transfer," *Cognitive Psychology*, 20, 362-404.
- Kreps, David W. and Robert Wilson (1982), "Sequential Equilibria," *Econometrica*, 50, 853-894.
- Kurland, D.M., Roy D. Pea, Catherine A. Clement, and Ronald Mawby (1986), *A Study of the Development of Programming Ability and Thinking Skills in High School Students*, Bank Street College of Education, Center for Children and Technology.
- Larrick, R. P., Morgan, J. N., and Nisbett, R. E. (1990), "Teaching the Use of Cost-Benefit Reasoning in Everyday Life," *Psychological Science*, 1, 362-370.
- Milgrom, Paul and John Roberts (1982), "Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis," *Econometrica*, 50, 443-459.
- Miller, Ross M. and Charles R. Plott (1985), "Product Quality Signalling in Experimental Markets," *Econometrica*, 53, 837-872.
- Mookherjee, Dilip and Barry Sopher (1997), "Learning and Decision Costs in Experimental Constant Sum Games," *Games and Economic Behavior*, 19, 97-132.
- Moulton, Brent (1986), "Random Group Effects and the Prevision of Regression Estimates," *Journal of Econometrics*, 32, 385-397.

- Nagel, Rosemarie (1995), "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85, 1313-1326.
- Pea, Roy D. and D.M. Kurland (1984), "On the Cognitive Effects of Learning Computer Programs," *New Ideas in Psychology*, 2, 137-168.
- Perkins, David N. and Gavriel Salomon (1988), "Teaching for Transfer," *Educational Leadership*, 46, 22-32.
- Perkins, David N., S. Schwartz and R. Simmons (1988), "Instructional Strategies for the Problems of Novice Programmers," in Richard E. Mayer, ed., *Teaching and Learning Computer Programming: Multiple Research Perspectives*, Hillsdale, NJ, Erlbaum, 153-178.
- Rankin, Frederick W., John Van Huyck, Ray Battalio (2000), "Strategic Similarity and Emergent Conventions: Evidence from Payoff Perturbed Stag Hunt Games," *Games and Economic Behavior*, 32(2), 315-337.
- Roth, Alvin, E. and Erev, Ido (1995), "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, 164-212.
- Roth, Alvin, Ido Erev, Robert Slonim and Greg Berman (2000), "Equilibrium and Learning in Economic Environments: The Predictive Value of Approximations," working paper.
- Salomon, Gavriel and David N. Perkins (1989), "Rocky Roads to Transfer: Rethinking Mechanisms of a Neglected Phenomenon," *Education Psychologist*, 24, 113-142.
- Stahl, Dale O. (1996), "Boundedly Rational Rule Learning in a Guessing Game," *Games and Economic Behavior*, 16, 303-330.
- Stahl, Dale O. (2000), "Action Reinforcement Learning versus Rule Learning," working paper.

Table 1a:

Monopolist Payoffs

High Cost Monopolist (MH)			Low Cost Monopolist (ML)		
Monopolist	Entrant Response		Monopolist	Entrant Response	
Action	IN	OUT	Action	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 1b:

Entrant Payoffs, High Cost Entrants

Entrant's Strategy	Monopolist's Type	
	High Cost	Low Cost
IN	300	74
OUT	250	250

Table 1c:

Entrant Payoffs, Low Cost Entrants

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

Table 2a:

Monopolist Payoffs, Price Game

Low Cost Monopolist (ML)		
Monopolist	Entrant Response	
Action	IN	OUT
1	204	545
2	333	678
3	355	700
4	378	723
5	350	695
6	283	648
7	250	615

High Cost Monopolist (MH)		
Monopolist	Entrant Response	
Action	IN	OUT
1	-428	-220
2	-298	-110
3	8	165
4	103	448
5	125	470
6	148	493
7	125	470

Table 2b:

Entrant Payoffs, Price Game, Low Cost Entrants

Entrant's Strategy	Monopolist's Type	
	Low Cost	High Cost
IN	219	594
OUT	281	281

Table 3
Summary of Experimental Treatments

	Control Sessions	Crossover Sessions Experiment 1	Crossover Sessions Experiment 2
Number of Sessions (Experienced Only)	5	5	3
Number of Subjects (Experienced Only)	65	66	38
Number of 12 Period Cycles, Inexperienced Sessions	3 cycles (except for one session that only had two cycles)	3	3
Number of 12 Period Cycles, Experienced Sessions	3 sessions w/ 4 cycles 2 sessions w/ 3 cycles	4	4
Twice-experienced subjects?	No	No	All the subjects in one session, half the subjects in a second session.
Timing of Crossover	---	After first cycle	After first cycle for 2 sessions, after second cycle for the third session ³²
Type of Entrants and Equilibrium (Inexperienced & Before Crossover)	Low Cost Entrants Separating Only	Low Cost Entrants Separating Only	High Cost Entrants Pooling and Separating
Type of Entrants and Equilibrium (Following Crossover)	---	Low Cost Entrants Separating Only	Low Cost Entrants Separating Only
Change in Presentation of Payoffs Following Crossover?	---	Flip from Quantity to Price (or vice versa)	None

³²The session with all once experienced subjects and all twice experienced subjects was crossed over after the 1st cycle. The session with half twice experienced subjects was crossed over after the 2nd cycle.

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** (.166)	1.960** (.111)	2.384** (.117)
Probability Change of Strategy (p_{change})	.503** (.023)	.645** (.029)	.674** (.026)
Discounting of Past Experience (δ)	 	.084** (.009)	.060** (.010)
Square Root of Weight on Initial Beliefs Following Crossover (ρ)	 	.054+ (.028)	.022 (.036)
Weight on Others' Experience (w_{Other}) (Multiplied by 100)	 	.259 (.162)	.699 (.454)
Probability Non-Learner	 	.310** (.053)	.254** (.046)
Probability Sophisticated Learner	 	 	.188** (.044)
Probability Switching Type Non-Learner to Unsophisticated	 	 	.073 (.072)
Probability Switching Type Unsophisticated to Sophisticated	 	 	.136** (.045)
Log Likelihood	-4517.11	-4215.31	-4120.60

Notes: The full data set has 4595 observations over 104 individuals, including 2585 observations from 66 individuals in the control sessions and 2010 observations from 38 individuals in the crossover sessions.

** statistically significant at the 1% level
+ statistically significant at the 10% level

*statistically significant at the 5% level

Appendix: Statistical Analysis of Experiments 1 and 2

We reached several conclusions in the text regarding the effects of the crossover treatments (namely Conclusions 2, 4 and the discussion of results in the last paragraph of section V). This appendix shows that these claims based on the raw data reported in the text hold up with formal econometric analysis of the data controlling for covariates affecting behavior. As in the text our focus here is on strategic play by MLs.

All of 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. (As in the main body of the paper, choices in the price game have been transposed to be in terms of quantities.) To correct for individual effects in the data, standard errors are calculated using Moulton's (1986) correction for clustering.³³ 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.

Experiment 1: Regressions on ML data from Experiment 1 are reported in Table A.1. The data set consists of all observations, both from inexperienced and experienced sessions, for all subjects who returned for an experienced subject session. Although the inexperienced sessions are of no direct interest, including this data permits better control of the individual effects in the data.

[Insert Table A.1. here]

Formally, the full specification for the latent variable underlying the ordered probit is given by the following equation: The variable O_t^i is the (latent) output for observation t of subject i . The variable CO is a dummy for the crossover treatment. The variables Cyc_{E2} , Cyc_{E3} , and Cyc_{E4} are dummies for the second, third, and fourth twelve period cycles of an experienced session. The variables Cyc_{I1} , Cyc_{I2} , and Cyc_{I3} are dummies for the first, second, and third twelve period cycles of an inexperienced session. 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 error term is given by ε_t^i .

³³See also White (1994). Rather than assuming that observations from the same individual are independent, this technique assumes that observations from the same individual are correlated and corrects the standard errors accordingly. This correction is a variation on the sandwich estimator and is the technique underlying the "cluster" option in Stata.

$$\begin{aligned}
O_t^i = & \alpha + \beta CO + \delta_1(Cyc_{E2} + Cyc_{E3} + Cyc_{E4}) + \delta_2(Cyc_{E3} + Cyc_{E4}) + \delta_3Cyc_{E4} + \delta_4(Cyc_{I1} + Cyc_{I2} + Cyc_{I3}) \\
& + \delta_5(Cyc_{I1} + Cyc_{I2}) + \delta_6(Cyc_{I1}) + \phi_1CO * (Cyc_{E2} + Cyc_{E3} + Cyc_{E4}) + \phi_2CO * (Cyc_{E3} + Cyc_{E4}) + \phi_3CO * Cyc_{E4} \\
& + \phi_4CO * (Cyc_{I1} + Cyc_{I2} + Cyc_{I3}) + \phi_5CO * (Cyc_{I1} + Cyc_{I2}) + \phi_6CO * (Cyc_{I1}) + \gamma ER + \eta Con + \epsilon_t^i
\end{aligned}$$

(Equation A.1)

Less formally, the independent variables used in the regressions can be broken into four general classes as follows:

1. *Controls for the Time Period:* The base here is the first twelve period cycle in the experienced subject sessions. Dummies are used to generate differences going both forward and backward in time from this base. Specifically, the regressions include the following time dummies: inexperienced sessions, periods ≤ 12 ; inexperienced sessions, periods ≤ 24 , inexperienced sessions, all periods; experienced sessions, periods ≥ 13 ; experienced sessions, periods ≥ 25 ; and experienced sessions, periods ≥ 37 . Because results for the inexperienced sessions are of no direct interest, all parameter estimates for the inexperienced sessions are suppressed on Table 1A. Because the time dummies have an overlapping structure, the associated parameter estimates capture differences between twelve period cycles. For example, a positive estimate of the coefficient for “Experienced, Periods ≥ 13 ” (δ_1 in equation A.1) tells us that output levels are higher in the second twelve period cycle of the experienced session than the first twelve period cycle of the experienced sessions.
2. *Controls for the Crossover Treatment:* These include a dummy for all subjects who eventually experienced a crossover (this dummy is coded as a 1 even before the subject actually goes through the crossover) as well as interactions between the crossover dummy and the time dummies. For purposes of interpretation, the crossover dummy gives the difference between the two treatments in the base time period (the first twelve period cycle of the experienced sessions) and the interactions then give differences in changes over time. For example, the parameter labeled “Crossover * Experienced, Period ≥ 13 ” (ϕ_1 in equation A.1) estimates the difference between the increase in output in the crossover treatment from the first to second twelve period cycles of the experienced sessions and the increase in output in the no crossover treatment between these two cycles. A negative parameter estimate for this variable doesn’t mean that output is decreasing in the crossover treatment (although this may also be true), but rather that output is increasing less than in the no crossover treatment. This “differences in differences” approach gives additional control beyond correcting the standard errors for clustering for the different base levels of strategic play by MLs in the two treatments.
3. *Controls for Es’ Behavior:* The control variables are the entry rates for the current twelve period cycle following outputs 2, 3, 4, 5, and 6. 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.³⁴ As a group, the entry rate controls are

³⁴Identical 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.

always easily significant at the 1% level. However, because these parameters are not of direct interest, they are suppressed in Table A.1. The second line of Table A.1 indicates whether these variables are included in the regressions.

4. *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. As these parameters are not of direct interest here they are also suppressed in Table A.1. Once again, the second line of the table indicates whether these control variables are included in a regression.

Turning to the regression results, Model 1 checks for a crossover effect without controlling for entry rates or context. The key parameter for measuring the crossover effect is the dummy “Crossover* Experienced Period ≥ 13 ” (ϕ_1 in the formal specification). This parameter estimate is negative and significant at the 1% level. Given the differences in differences structure of the econometric model, this result indicates that output is increasing less between the first and second twelve period cycles in sessions with a crossover than those without. This regression confirms that the stalled learning observed in Figure 2 is significant even after controlling for individual effects in the data. The insignificant parameter estimates for “Crossover, Experienced, Period ≥ 25 ” and “Crossover, Experienced, Period ≥ 37 ” (ϕ_2 and ϕ_3 in the formal specification) indicating that the growth in strategic play by MLs in the cross-over treatment does not catch up to the controls over the last two 12 period cycles.

This negative parameter estimate for “Crossover*Experienced Period ≥ 13 ” in Model 1 does not imply a drop in output for the crossover sessions. Adding together the dummies for “Experienced, Period ≥ 13 ” and “Crossover*Experienced, Period ≥ 13 ,” the total change for crossover sessions is very slightly positive and not significantly different from zero. This confirms that the crossover puts a temporary halt to increases in strategic play, rather than reducing it, and definitely does not return experienced MLs to the level of inexperienced subject play.

Model 2 adds the five entry rate controls to the regression and Model 3 adds to this the controls for use of meaningful context. The control variables for entry rates easily achieve joint statistical significance at the 1% level while the controls for use of meaningful context are just below the cutoff for significance at the 5% level. Adding these controls lessens the magnitude of the crossover effect as well as reducing its statistical significance. However, in both models, the parameter estimate for “Crossover*Experienced, Period ≥ 13 ” still achieves statistical significance at the 5% level. In neither model do the parameter estimates for “Crossover” (β in equation A.1) achieve statistical significance at any standard level, indicating no significant differences between the two treatments for experienced subjects prior to the cross-over.³⁵ Further, both the “Crossover*Experience, Period ≥ 25 ” and the “Crossover*Experience, Period ≥ 37 ” parameters fail to achieve statistical significance at standard levels, confirming the absence of any catchup in the cross-over treatment over the last two cycles.

Experiment 2: Regressions on ML data from Experiment 2 are reported in Table A.2. The data set for these regressions includes all data from games with low cost Es for subjects who returned for an experienced subject session. Note that data from games with high cost Es (data prior to the crossover) are not included.

³⁵In the results section for Experiment 1, we assert that there is no significant difference between play in the price and quantity games prior to the crossover. To test this proposition formally, we ran ordered probits using all choices by MLs prior to the crossover in Experiment 1. The right hand side variables were dummies for the three twelve period cycles of the inexperienced sessions, a dummy for the first twelve period cycle of the experienced session, and interactions of these variables with a dummy for use of the price game. (The dummy for the first cycle of the inexperienced session is dropped to avoid colinearity.) While the sign of the price game dummies is consistently positive, none of these dummies was statistically significant even at the 10% level. Thus, strategic play was slightly more prevalent in the price game, but not to any significant degree.

[Insert Table A.2 here]

Formally, the full specification for the latent variable underlying the ordered probit is given by equation A.2 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 Cyc₁₂ and Cyc₁₃ 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. Note that the time dummies do not have an overlapping structure in this specification – we are measuring levels, not differences. 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 .

$$O_t^i = \alpha + \beta_1 Lo * Cyc_{12} + \beta_2 Lo * Cyc_{13} + \beta_3 Lo * Cyc_{E1} + \beta_4 Lo * Cyc_{E2} + \beta_5 Lo * Cyc_{E3} + \beta_6 Lo * Cyc_{E4} + \delta_1 Hi * Cyc_{CR1} + \delta_2 Hi * Cyc_{CR2} + \delta_3 Hi * Cyc_{CR3} + \gamma ER + \eta Con + \lambda_1 SMH + \lambda_2 SML + \tau TCRS + uIMT + \epsilon_t^i$$

(Equation A.2)

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. Unlike the regressions for Experiment 1, the parameter estimates on the time dummies represent levels, not differences.³⁶ Thus, a positive parameter estimate for “crossover, periods 25-36 after crossover” tells us that there is more strategic play by MLs in this time period than in the first twelve period cycle of play by inexperienced subjects in the control sessions. Parameter estimates for the inexperienced sessions are suppressed in Table A.2 since they are not of any direct interest.
2. *Controls for Es’ Behavior:* These are identical to the controls used in the analysis of experiment 1. These parameter estimates are also suppressed in Table A.2 since they are not directly relevant. The second line of Table A.2 indicates whether these variables have been included in a model.
3. *Controls for Meaningful Context:* These are identical to the controls used in the analysis of experiment 1. The parameter estimates for these variables are not reported in Table A.2, 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

³⁶The natural reference point in Experiment 2 is the (absolute) level of strategic play in the cross-over treatment compared to inexperienced controls. In Experiment 1 the natural reference point is the change in the growth rate of strategic play compared to the experienced controls following the cross-over.

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 in Table A.2 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 at the 1% level.³⁷ 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 -.028 with a standard error of .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 individually³⁸ and the three crossover dummies fail to be jointly significant ($\chi^2 = 0.29$, 3 d.f., $p > .10$). 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 ($\chi^2 = 70.08$, 12 d.f., $p < .01$), 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.³⁹ While neither the timing of the crossover (τ in equation A.2) nor direct experience with an ML playing strategically (ι in equation 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

³⁷The parameter estimates are 1.022 and 1.012 respectively with standard errors of .239 and .229.

³⁸The parameter estimates are -.043 and -.046 respectively with standard errors of .263 and .224.

³⁹We 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.

equation A.2) are positive and statistically significant at the 1% level.

Table A.1

Experiment 1, Crossover Effects for MLs
 Ordered Probits, Standard Errors Corrected for Clustering
 Dependent Variable: Output Level

	Model 1	Model 2	Model 3
Control Variables	None	Entry Rates	Entry Rates Context
Number of Parameters	19	24	31
(δ_1) Experienced Period ≥ 13	.473** (.124)	.401** (.135)	.577** (.170)
(δ_2) Experienced Period ≥ 25	.220* (.104)	.195+ (.102)	.103 (.136)
(δ_3) Experienced Period ≥ 37	.206 (.153)	.241 (.171)	.175 (.169)
(β) Crossover	.311+ (.172)	.286 (.183)	.281 (.191)
(ϕ_1) Crossover * Experienced Period ≥ 13	-.469** (.171)	-.385* (.185)	-.385* (.189)
(ϕ_2) Crossover * Experienced Period ≥ 25	.066 (.152)	.083 (.156)	.089 (.156)
(ϕ_3) Crossover * Experienced Period ≥ 37	.181 (.197)	.054 (.213)	.045 (.212)
Log Likelihood	-3271.26	-3256.91	-3249.88

Notes: All regressions contain 2696 observations over 131 individuals. Only individuals who returned for an experienced session are included. No observations from twice-experienced subjects are included.

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

Table A.2

Experiment 2, Crossover Effects for MLs: Ordered Probits,
Standard Errors Corrected for Clustering, Dependent Variable is Output Level

	Model 1	Model 2	Model 3	Model 4
Control Variables	None	Entry Rates	Entry Rates Context	None
Number of Parameters	15	20	27	19
(β_3) No Crossover, Experienced Periods 1 - 12	.613** (.151)	.750** (.213)	.395+ (.226)	.640** (.158)
(β_4) No Crossover, Experienced Periods 13 - 24	1.095** (.155)	1.268** (.198)	.902** (.197)	1.145** (.163)
(β_5) No Crossover, Experienced Periods 25 - 36	1.323** (.161)	1.418** (.176)	1.053** (.201)	1.384** (.170)
(β_6) No Crossover, Experienced Periods 37 - 48	1.534** (.191)	1.560** (.213)	1.291** (.289)	1.604** (.204)
(δ_1) Crossover Periods 1 - 12 After Crossover	.585** (.151)	.592** (.167)	.436* (.175)	.529** (.194)
(δ_2) Crossover Periods 13 - 24 After Crossover	1.052** (.240)	1.055** (.241)	.915** (.253)	1.010** (.200)
(δ_3) Crossover Periods 25 - 36 After Crossover	1.277** (.206)	1.355** (.212)	1.238** (.222)	1.380** (.256)
(λ_1) Strategic Play as MH 10 Plays Prior to Crossover				.123** (.025)
(λ_2) Strategic Play as ML 10 Plays Prior to Crossover				.235** (.056)
(τ) Cycle When Crossover Occurs				.022 (.089)
(ι) Paired with an ML Who Played Strategically				.238 (.193)
Log Likelihood	-1999.67	-1993.31	-1964.63	-1929.74

Notes: All regressions contain 1654 observations over 104 individuals. Individuals who were crossed from the game with high cost entrants (Table 1b) to the game with low cost entrants (Table 1c) are included as well as individuals from the no crossover cell of Experiment 1.

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