Lessons Learned: Generalizing Learning Across Games

David J. Cooper and John H. Kagel*

This paper synthesizes findings from an ongoing research program on learning in signaling games. The present paper focuses on cross-game learning - the ability of subjects to take what has been learned in one game and generalize it to related games - an issue that has been ignored in most of the learning literature. We begin by laying out the basic experimental design and recapitulating early results characterizing the learning process. We then report results from an initial experiment in which we find a surprising degree of positive cross-game learning, contrary to the predictions of commonly employed learning models and to the findings of cognitive psychologists. We next explore two features of the environment that help to explain when and why this positive transfer occurs. First, we examine the effects of abstract versus meaningful context, an issue that has been largely ignored by economists out of the belief that behavior is largely dictated by the deep mathematical structure of a game. In contrast, results from cognitive psychology suggest that behavior may well be sensitive to context employed. Our results show that the use of meaningful context serves as a catalyst for positive transfer. Second, we explore how play by two-person teams differs from play by individuals. The psychology literature is quite pessimistic about the ability of teams to beat a "truth wins" standard based on performance of individuals. But teams easily surpass this norm in our cross-game experiment. We use the dialogues between team members to gain insight into how this transfer occurs, gaining direct confirmation for hypotheses generated by econometric analysis of earlier data.

I. The Experimental Environment: Our experiments are based on a simplified version of Paul Milgrom and John Roberts' (1982) entry limit pricing game. The game proceeds as follows: (1) Monopolists (Ms) observe their cost level - high (MH) or low (ML) cost - realized according to equal probabilities that are common knowledge. (2) Ms choose a quantity (output) whose payoff is contingent on the entrant's (Es) response (see Table 1). (3) E sees this output, but not M's type, and either enters or stays out. 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).

[Insert Table 1 here]

As a treatment variable, Es use either the high or low cost payoff table. With high cost Es, there exist pooling (sequential) equilibria at outputs 2 - 5, as well as two separating equilibria in which MHs choose 2 and MLs choose either 6 or 7. Standard equilibrium refinements leave only the pooling equilibria at output 4 or 5 and the separating equilibrium with MLs choosing 6. In the game with low cost Es, no pooling equilibria exist but the two separating equilibria remain.

Experiments employ between 12-16 subjects with half of them playing as Ms, the other half playing as Es. A round robin format is used where each M meets a different E for several plays of the game. Es and Ms then switch roles and play another set of games. This process repeats itself for a total of between 24-48 games.

Following each play of the game Es are informed of M's type. In addition subjects are informed of choices for *all* Ms, the Es they were paired with, and Ms type so that they can readily form population expectations for Ms' choices and Es' responses.

Games are characterized using either abstract or meaningful context. With abstract context neutral terms are used throughout; e.g., Ms are called "A players" (type A1 for MH, type A2 for ML) and Es are called "B players." The instructions provide subjects with no guidance as to the situation being modeled. With meaningful context natural (but nonpejorative) terms are used; e.g., Ms are called "existing firms" ("high cost" and "low cost" types) and Es are called "other firms" deciding between "entry into this industry or some other industry." Although we would be surprised if our subjects had direct experience with this context, it does provide a meaningful framework for organizing their thoughts about the game.

Some sessions have employed subjects playing in two person teams throughout. Team members coordinate their choices through an instant messaging system, with these messages recorded for content analysis, thereby providing a window into the learning process.

Early experiments reveal a consistent dynamic to the learning process (David J. Cooper, Susan Garvin and John H. Kagel, 1997a, b): In early periods, both types of Ms typically ignore the strategic possibilities of the game, choosing their "myopic maxima" – the output that maximizes payoffs *ignoring* the threat of entry (2 for MHs, 4 for MLs) – even though entry rates make strategic play incentive compatible from the start. Reacting to high entry rates at output 2, MHs gradually learn to play strategically by choosing 4. In games with high cost Es, play converges to the pooling equilibrium at 4. In games with low cost Es, no such equilibrium exists. Instead, as entry rates for output

4 rise, MLs gradually learn to play strategically by choosing outputs 5 or 6. This process is very slow – even by the end of inexperienced subject sessions MLs' modal choice is 4. Only in experienced subject sessions does strategic play become the modal choice for MLs. Intuitively, playing strategically is more difficult to learn as an ML in the game with low cost Es than as an MH in the game with high cost Es because MHs can rely on imitating the choices of MLs, while MLs have no such guide to follow in early rounds.

The evolution of Ms' choices in the limit pricing game is consistent with a simple model of adaptive learning based on fictitious play (Cooper, Garvin and Kagel, 1997 a, b). We don't claim that this adaptive learning model is the "best" model of learning (if any "best" model of learning exists). Rather, we find that models based on fictitious play do a good job of tracking data from our experiments, generating new hypotheses that can be tested experimentally, and providing a useful framework for organizing results from new experiments.

II. An Initial Experiment with Positive Transfer: In our initial cross-game learning experiment the most striking results come from a treatment where subjects first play a full session of the game with high cost Es and, after several plays of this game as experienced subjects, are switched to the game with low cost Es (Cooper and Kagel, 2001). As experienced subjects, prior to the crossover, play has converged strongly to the pooling equilibrium at 4.

The pooling equilibrium requires MHs to play strategically, with no strategic play on the part of MLs. However, following the crossover, equilibrium requires that MHs not play strategically, while MLs must play strategically. Given that virtually all subjects have learned to play strategically as MHs prior to the crossover, the critical question is whether this experience will help them to play more strategically following the crossover as MLs than inexperienced subjects (positive transfer).

The large psychology literature dealing with learning generalizability provides little reason to be optimistic regarding high levels of positive transfer (see, for example, Gavriel Solomon and David N. Perkins, 1989). Moreover, our adaptive learning model predicts, if anything, *negative* transfer. Intuitively, fictitious play learners treat their opponents like a fixed statistical distribution, rather than forming a model of how their opponents make decisions. When the payoffs change, they do not immediately anticipate a change in their opponents' behavior. Instead, they only begin to play strategically as MLs when their beliefs from prior to the crossover have been unlearned. Any learning model that does not allow players to anticipate that a change in their opponents' payoffs will lead to a change in their behavior will yield similar predictions.

Contrary to these predictions, MLs showed significantly more strategic play immediately following the crossover than in inexperienced control sessions; 25.7% in the first cycle following the crossover versus 8.5% in the first cycle of the control sessions (see Control 1 and Cross 1 in Table 2).[1] In fact, strategic play by MLs following the crossover is statistically indistinguishable from experienced subject play in control sessions, suggesting that experience with the pooling equilibrium in the high cost entrant game is an almost perfect substitute for experience in the low cost entrant game.[2]

To get the adaptive learning model to track the data, we add sophisticated learners who model their opponents as unsophisticated learners and thereby anticipate the increase in entry rates following the crossover. 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 with experience*. It is the increase in the proportion of sophisticated learners that generates the positive cross-game transfer.

This initial experiment raises two questions. First, what features of the environment encourage (or discourage) positive transfer? Second, our explanation for positive transfer rests on statistical estimates of a structural learning model. Can we find direct evidence that sophistication increases as a result of past experience and results in positive transfer?

III. Meaningful Context: This initial cross-game learning experiment employed meaningful context. We speculated that this use of meaningful context might have played a role in generating the positive transfer. The psychology literature reports abundant evidence that subjects' ability to solve logic problems can be greatly improved by the use of meaningful (as opposed to abstract) context (see, for example, Roger L. Dominowski, 1995). To the extent that learning to play strategically is similar to solving a logic problem, we might expect that meaningful context will facilitate the development of strategic play. Indeed, Cooper and Kagel (2002) found that meaningful context facilitated the development of strategic play (particularly for MHs) in limit pricing games which did not involve any cross-over treatments.

Cooper and Kagel (2003) replicate this cross-over experiment using abstract context. In this case we find that strategic play of MLs following the cross-over is less than for inexperienced controls; 7.1% versus 15.1% for the first cycle (see Table 2, Control 2 and Cross 2). This difference, while small, is statistically significant even controlling for Es' choices. The contrast with the meaningful context is even more striking since the increase in entry rates for output 4 following the crossover make strategic play by MLs highly incentive compatible in both experiments.[3] These results provide clear evidence that meaningful context can facilitate positive transfer between related games. Given the hypothesized relationship between positive transfer and the presence of sophisticated learners, these results also suggest that meaningful context not only speeds up learning as reported in Cooper and Kagel (2002), but can actually change the nature of subjects' reasoning process.

IV. Team play: As an alternative catalyst for positive transfer we have explored play by two-person teams (Cooper and Kagel, 2003). The motivation for this treatment is straightforward: Many economic decisions are made within a team framework. As a result, if there are major differences between individual and team outcomes, extrapolation of individual outcomes will be inconsistent with observed behavior in many settings. Further, by having subjects play in teams and recording their discussions, we hoped to obtain direct evidence about the process underlying positive transfer.

These experiments used abstract context throughout. However, unlike the abstract context treatment with individual subjects, there is substantial positive transfer (Table 2, Control 3 and Cross 3), far stronger than was observed for the meaningful context treatment with individuals. To truly appreciate the effectiveness of teams, consider a demanding norm from the psychology literature known as "truth wins." For Ms, learning to play strategically is similar to solving a puzzle, the solution to which, once discovered, is likely to be self-confirming. In this setting, absent any synergies from interactions between team members, team performance should reflect the insights of the most strategic member of the team; i.e., if an individual subject has figured out the benefits of strategic play with probability p, the probability of a team solution, P, is $P = 1 - (1-p)^2$. This demanding standard is rarely exceeded, sometimes matched, and usually not achieved in psychology experiments (see James H. Davis, 1992, for a review of the literature). (Psychologists attribute this relative "inefficiency" of groups to reduced member motivation and/or coordination problems in combining team members contributions, both of which are familiar problems in the economics literature on team production.) The increased level of strategic play for MLs following the crossover easily exceeds the "truth wins" standard (based on individuals in the abstract context treatment). Team play seems to fundamentally change how subjects play following the crossover.

The dialogues between team members provide considerable insight into the mechanism underlying positive transfer. First, almost all teams realize that the change in Es payoffs will increase entry rates for currently chosen outputs – 76% are coded as recognizing this in the first play of the game following the cross-over - regardless of whether they were playing as Ms or Es. For example, from one M team, "Now they are making this interesting. ...Now it'd actually be worth it to pick [IN] every time." Second, 45% of teams in the first cycle of play following the crossover discussed how higher outputs as MLs would make Es recognize that they weren't MHs. For example, "If we enter 6 when we are [an ML] then everyone will know that we are an [ML] and will guess accordingly, giving us a higher average....The 6 pays less but is very clear what we are." In comparison, less than 18% of teams in the first cycle of play in control sessions are coded for such comments. To summarize, analysis of the dialogues indicates that teams anticipate that changes in Es' payoffs will lead to changes in Es' behavior, and the proportion of teams discussing strategic play by putting themselves in the shoes of Es is increased by experience in games with high cost Es. This closely mirrors the explanation for positive transfer proposed in Cooper and Kagel (2001).

V: Summary and conclusions: Recent experiments with limit pricing games find surprising levels of positive transfer between games with very different equilibria. Econometric analysis of the experimental data and analysis of dialogues between teammates suggest that this positive transfer relies on (i) the existence of sophisticated learners who correctly anticipate the impact of changes in their opponents' payoffs on their opponents' choices and (ii) the proportion of sophisticated learners growing with experience, including experience in related games. This growth in sophisticated learners, with the resulting positive transfer between games, is facilitated by the use of meaningful context or the use of two-person teams.

Our results contain surprises for both economists and psychologists. The impact of meaningful context is not something that standard economic theory, with its emphasis on procedural invariance would anticipate, nor can economic theory readily accommodate such context effects. Two of our results -- positive transfer and team performance beating the "truth wins" criterion -- were unanticipated on the basis of the psychology literature. The differences between what we have found and typical psychology experiments cannot be attributed to our use of meaningful context, as psychologists' studies of learning generalizability invariably use meaningful context. There is, however, one important difference between our experiments and those psychologists typically perform: Psychologists tend to study individual learning in one-shot environments. Our games involve repeated interactions between agents. These repeated interactions appear to be responsible for both these effects as (i) they underlie the development of increased numbers of sophisticated learners that accounts for the positive cross-game learning and (ii) comparisons of teams with individual controls show that it takes several plays of the game before teams are able to beat the "truth wins" norm (Cooper and Kagel, 2003).

*Department of Economics, Case Western Reserve University, Cleveland, OH 44106 and Department of Economics, Ohio-State University, Columbus, OH 43210, respectively. Research support from the National Science Foundation is gratefully acknowledged.

1.A cycle refers to a set of 12 (or 8) games in which subjects played the first 6 (or 4) games as Es or Ms and then switched roles for the remaining games.

2.Both these results continue to hold after controlling for any differences in entry rates between the cross-over treatment and the controls.

3.These abstract context sessions were conducted at Ohio State, the meaningful context sessions at the University of Pittsburgh. There were also a number of changes in the software for the abstract sessions, designed to match as closely as possible the team sessions. These differences may explain the substantially higher levels of strategic play between Controls 1 and 2. To ensure that the cross-over results we are attributing to context do not result form these differences we are replicating the meaningful context sessions at Ohio State using the same software (but with meaningful context) as the abstract context sessions. While these sessions haven't been completed, preliminary results show positive cross-game transfer for individual subjects with meaningful context.

References

Cooper, David J., Garvin, Susan and Kagel, John H. "Signaling and Adaptive Learning in an Entry Limit Pricing Game." RAND Journal of Economics, Winter 1997, 28 (4), pp. 662-83 (a).

_______, _______, ______. "Adaptive Learning vs. Equilibrium-Refinements in an Entry Limit Pricing Game." The Economic Journal, May1997, 107 (442), pp. 553-75 (b).

________ and Kagel, John H. "Learning and Transfer in Signaling Games." Working paper, Case Western Reserve University, 2001.

_______ and ________. "The Impact of Meaningful Context on Strategic Play in Signalling Games." Journal of Economic Behavior and Organization, 2002, (in press).

_________ and _______. "Are Two Head Better than One? Team versus Individual Play in Signaling Games." Working paper, Case Western Reserve University/Ohio State University, 2003.

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

Dominowski, Roger L. "Content Effects in Wason's Selection Task" in Stephen Newstead and Johnanthan St. B. T. Evans, eds., Perspectives on thinking and reasoning: Essays in honor of Peter Wason. Hillsdale, Erlbaum Associates, 1995, pp. 41-67.

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

Solomon, Gavriel and Perkins, David N. "Rocky Roads to Transfer: Rethinking Mechanisms of a Neglected Phenomenon." Education Psychologist, 1989, 24(2), pp. 113-42.

Table 1

Monopolist Payoffs

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

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

Entrant Payoffs

Entrant's Strategy	Monopolist's Type	
	High Cost	Low Cost
IN	300 (500)	74 (200)
OUT	250	250

Table 1: Payoff tables: Payoffs for high and low cost Es for OUT are the same. Low cost Es' payoffs for IN are in parentheses.

Table 2

Percentage of Strategic Play by MLs

Crossover vs. Control Sessions

	Control Sessions (No Crossover)				Crossover Sessions		
	Control 1	Control 2	Control 3		Cross 1	Cross 2	Cross 3
Meaningful Context	U			Meaningful Context	U		
Teams			U	Teams			U
				Cycle Preceding Crossover	.052	.000	.050
Inexperienced Cycle 1	.085	.151	.273	Post-Crossover Cycle 1	.257	.071	.536
Inexperienced Cycle 2	.136	.406	.624	Post-Crossover Cycle 2	.469	.245	.777
Inexperienced Cycle 3	.234	.490*	.754**	Post-Crossover Cycle 3	.535	.550	.902

* Only half of the inexperienced sessions included a third cycle.

** Inexperienced team sessions only included two cycles. Data reported here is from the first cycle of experienced play.