

Learning from Peers in Signaling Game Experiments*

Xiaodong Liu^{a†}, John H. Kagel^b and Lung-fei Lee^b

^aDepartment of Economics, University of Colorado, Boulder, CO 80309, USA

^bDepartment of Economics, Ohio State University, Columbus, OH 43210, USA

March 30, 2011

Abstract

We investigate peer group effects in laboratory experiments based on Milgrom and Roberts' (1982) entry limit pricing game. We generalize Heckman's (1981) dynamic discrete choice panel data models by introducing time-lagged social interactions, using the unbiased GHK simulator to implement the computationally cumbersome maximum likelihood estimation. We find that subjects' decisions are significantly influenced by past decisions of peers on several dimensions including potential entrants' choices and strategic play of like-type monopolists. The proposed model and estimation method may be applicable to other experiments where peer group effects are likely to play an important role.

Key Words: Learning, Peer Group Effects, Signaling Game Experiments, Discrete Choice, Dynamic Models, Simulated Maximum Likelihood.

JEL Classification: C15, C35, C72, C90

*We appreciate valuable comments and suggestions from the co-editor Edward Vytlačil and three anonymous referees to improve the presentation of this paper. Kagel appreciates financial support from NSF under grant numbers SES-0136928 and SES-0452911. Lee appreciates financial support from NSF under grant number SES-0519204. Any opinions, findings, and conclusions or recommendations in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. An earlier version of this paper circulated under the title "Dynamic discrete choice models with lagged social interactions: with an application to a signaling game experiment" was presented at the 2006 Far Eastern Meeting of the Econometric Society, Beijing, China, and the 2007 AEA Annual Meeting, Chicago. We thank participants at meetings for valuable comments. We are responsible for any errors.

†Correspondence to: Xiaodong Liu, Department of Economics, University of Colorado at Boulder, UCB 256, Boulder, Colorado, 80309, USA. E-mail: xiaodong.liu@colorado.edu. Phone: 1(303)492-7414. Fax: 1(303)492-8960.

1 Introduction

Many experiments observe that subjects initially tend to ignore the strategic implications of their actions and only gradually learn to make decisions that are consistent with equilibrium predictions. Further, the typical experimental design has a group of subjects interacting with each other over a number of replications of the same, or closely related, games. In cases where these interactions involve random rematching of players, there is some concern for “session level effects” as subjects might be influenced by peers’ choices. These potential session level effects are akin to what econometricians call “peer group” or “neighborhood” effects.

The present paper studies peer group effects in a signaling game experiment. Signaling game experiments are particularly susceptible to peer group effects, as ever since the pioneering work of Miller and Plott (1985), feedback to players typically includes population data — the signals *all* strong and weak types send and responses to same — as otherwise development of strategic play can be painfully slow. As such peer group effects are likely to be particularly pronounced.

To investigate the role of peer group effects on the development of strategic play, we generalize Heckman’s (1981) dynamic discrete choice panel data model by introducing time-lagged social interactions, so that the model can accommodate the interrelationships between decisions of cross-sectional units (peers). Likelihood functions of dynamic discrete choice models involve multiple integrals in cases where explanatory variables include lagged latent dependent variables or disturbances allow for serial correlation. For panel data models, the dimension of integration increases with the number of periods, which makes numerical implementation of maximum likelihood estimation impractical. To overcome this computational difficulty, simulation-based estimation methods developed by Geweke (1991), Borsch-Supan and Hajivassiliou (1993) and Keane (1994) have been shown to be practical and accurate for estimation purposes when time periods are not too long. We show that the implementation of the Geweke-Hajivassiliou-Keane (GHK) simulator remains tractable for the generalized model with social interactions. These procedures for estimating peer group effects are potentially applicable to a variety of experimental designs where subjects receive feedback on peers’ performance.

We apply the generalized model to study the evolution of a monopolist’s strategic play in experiments based on Milgrom and Roberts’ (1982) entry limit pricing game. Entrants’ responses to monopolists’ past choices are an obvious peer group effect in games of this sort. In what follows

we distinguish between the weight monopolists place on entrants responses to their own choices as opposed to responses to peers' choices. In addition we explicitly incorporate the influence of peer monopolists' choices on a monopolist's own choice.¹ We find that both types of peer group effects are relatively unimportant in games that have a pure strategy pooling equilibrium (which play reliably converges to when it exists). Further, what peer group effects there are, are primarily associated with a monopolist's responses to like-type monopolists' choices as opposed to potential entrants' responses. However, in games where no pure strategy pooling equilibrium exists, peer group effects are quite prominent as low cost monopolists learn to distinguish themselves from high cost types. In this case, peer group effects are associated both with like-type monopolists' choices as well as entrants' responses to these choices. Further, for experienced subjects at least, considerably more weight is placed on the behavior of the entrant a monopolist is paired with than the more informative subject population data.

From the statistical inference point of view, costs are likely to limit the number of experimental subjects, rounds, and sessions for any given experimental investigation, which might prevent one from determining whether peer group effects are negligible or are overwhelmed by estimation errors caused by insufficient sample size. As such, the current work sheds some light on sample size requirements and data structures favorable to successful identification of potential peer group effects in discrete choice games.

The paper proceeds as follows. Section 2 presents the equilibrium predictions of the entry limit pricing game, outlines the experimental procedures and provides a general description of the experimental data. It also provides the motivation for the empirical studies. In Section 3, we introduce a general dynamic discrete choice panel data model with lagged social interactions, derive the likelihood function and illustrate the formulation of simulators and simulated likelihood function for this model. Section 4 develops empirical econometric models to investigate the adjustment process of subjects's behavior in the entry limit pricing game and interprets the estimation results. Section 5 briefly concludes.

¹Past work with these data sets has explicitly incorporated entrants' responses to monopolists' choices. However, past work has not distinguished between entrants' responses to own choices relative to peers' choices, or explicitly incorporated the effects of like-type monopolists' choices on own choices. The latter has been dealt with through introducing clustering at the session level or through session level random effects.

2 The Signaling Game Experiments

2.1 Theoretical Predictions

The experiment is based on a simplified version of Milgrom and Roberts' (1982) entry limit pricing game. In the model there are two firms, an established monopolist M and a potential entrant E , in a two-stage game producing a homogeneous good. Nature decides M 's cost of production along with the distribution of these costs. M 's cost is private information throughout the game, with the prior distribution of costs being common knowledge. In the first stage, M observes his own cost and chooses an output level. In the second stage, E chooses to enter or stay out in response to the observed output level. The predetermined opportunity cost to E for entering the market is common knowledge. If entry occurs, Cournot duopoly profits are realized for both M and E . If there is no entry, M receives the single period monopoly profit. Entry is profitable against high cost M s (M_{HS}) but not against low cost M s (M_{LS}). There are two types of pure strategy equilibria: *pooling* equilibria in which M_{HS} produce at higher output levels than the single period profit maximizing level in order to make entry appear unattractive and *separating* equilibria in which M_{LS} produce at higher output levels for the same reason.

In this game, the information sets are defined by the realized costs of M and E (c_M and c_E) and M 's output level, Q . A (pure) strategy for M is a mapping s from possible cost levels into possible choices of Q and a (pure) strategy for E is a mapping t from R^2 into $\{0, 1\}$ for each possible pair (c_E, Q) , where 1 is interpreted as "enter" and 0 as "stay out". An equilibrium consists of a pair of strategies (s^*, t^*) and a pair of conjectures (\bar{s}, \bar{t}) such that (i) M 's pricing policy s^* is a best response to its conjecture \bar{t} about E 's entry rule, (ii) the strategy t^* is a best response for E to its conjecture \bar{s} , and (iii) the actual and conjectured strategies coincide. With two cost levels (types) for M , namely, $\underline{c}_M < \bar{c}_M$, $s^*(\underline{c}_M) = s^*(\bar{c}_M)$ is a *pooling* equilibrium; and $s^*(\underline{c}_M) \neq s^*(\bar{c}_M)$ is a *separating* equilibrium. Partial pooling is a mixed-strategy equilibrium with $s^*(\underline{c}_M) = s^*(\bar{c}_M)$ with a positive probability (less than 1). In a pooling equilibrium, E infers nothing about M 's type after observing Q and enters if the expected profit is positive. In a pure-strategy separating equilibrium, observing Q allows E to infer c_M exactly. Depending on M 's and E 's costs, the distribution of types, and the market demand function, pooling and/or separating equilibria can occur (Milgrom and Roberts, 1982, pp. 446-448).

[Tables 1-3 approximately here]

The entry limit pricing game is investigated experimentally in Cooper et al. (1997a; 1997b) and Cooper and Kagel (2003a; 2003b; 2004). In the experiments, the game is further simplified by adding the payoffs of the two stages together and providing subjects with payoff tables 1-3. In each round of the experiment, M is either a high-cost type (M_H) or a low-cost type (M_L) with equal probability, and chooses over output levels 1-7 in payoff table 1.² E 's cost is common knowledge. In any given experiment, E s are either all high cost types (E_{HS} ; payoff table 2) or all low cost types (E_{LS} ; payoff table 3) for all rounds. With E_{HS} , there exist pure-strategy pooling equilibria at output levels 1-5. There also exist two pure-strategy separating equilibria, in which M_H s always choose 2 and are always entered on, M_L s always choose 6 or 7 and are never entered on. Among them, only pooling at 4 or 5, and separating with M_L s choosing 6 survive Cho-Kreps' (1987) intuitive criteria.³ With E_{LS} , no pure-strategy pooling equilibrium exists, while the two pure-strategy separating equilibria still exist. There also exist a number of mixed-strategy equilibria. One that is of particular relevance is the partial pooling equilibrium in which M_L s always choose 5 while M_H s mix between 2 (with probability 0.8) and 5 (with probability 0.2), and E s always enter on choices other than 5, entering on 5 with probability 0.11. In simulations using a stochastic fictitious play learning model, this partial pooling equilibrium emerges with high frequency in the presence of E_{LS} (Cooper et al., 1997b). Further, in practice M_L s choose 5 with relatively high frequency as a separating equilibrium emerges (especially early on) and there is relatively little entry in response to it (Cooper et al., 1997b).

2.2 Experimental Procedures and Data

Detailed description of the experimental procedures can be found in Cooper et al. (1997b). The following lists some elements that are especially noteworthy in the statistical analysis that follows.

1. Each experimental session employed between 12 and 16 subjects who were randomly assigned to computer terminals. Sessions typically lasted 36 periods, with the number of periods announced in advance. Subjects switched roles after every six periods, with M s becoming E s

²Some of the experimental sessions reported employ a linear transformation of payoff tables 1 and 3, designed to represent a "pricing" version of the game (Cooper and Kagel, 2004). The two versions of the game are theoretically identical with the data pooled in the analysis after appropriate transformation of nominal choice values.

³The multiplicity of equilibria in games of this sort rests on the fact that appropriate out-of-equilibrium beliefs can serve to justify almost any outcome. For example, pooling at output level 1 can be justified by the belief that any deviation from 1 is more likely to result from an M_H than an M_L . This belief seems strange to say the least. Equilibrium refinements, one of the most accepted of which is Cho-Kreps intuitive criteria, are designed to rule out these strange beliefs on the basis of careful forethought.

and vice versa. We refer to a block of 12 periods with subjects playing the role of both M and E six times as an experimental cycle. In each period, M s' type is generated randomly.

2. Following each play of the game the outcomes from all pairings (M s' choice, E s' choice, and M s' type) were revealed to all subjects. This provides the basis for the potential peer group effects studied here.
3. Subjects were randomly paired with each other for each play of the game, with subject identification numbers suppressed. Hence there was no opportunity for reputation effects to develop. Learning, to the extent that it occurred, had to be based on own experience and observations of peers' choices and outcomes.

[Table 4 approximately here]

Experimental treatments are summarized in Table 4. Experienced subject treatments recruited subjects who had participated in earlier experimental sessions with exactly the same payoff tables. The treatment "Meaningful Context" uses natural language for the instructions, and was introduced to explore the effects of context on subjects' reasoning processes in signaling games (Cooper and Kagel, 2003a).

Subjects were recruited through announcements in undergraduate classes, posters placed throughout campus, advertisements in the campus newspapers, and direct e-mail contact. This resulted in recruiting a broad cross section of mostly undergraduate students and some graduate students. 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. Inexperienced-subject sessions lasted approximately one and a half hours, with experienced-subject sessions substantially shorter than this as short, summary instructions were used and subjects were familiar with the game. Subjects were paid between \$5-\$6 for showing up on time with total earnings averaging between \$26-\$27 per subject.⁴

2.3 Learning and Peer Group Effects

According to payoff table 1, with full information, choices of 2 and 4 are optimal for M_{HS} and M_{LS} respectively. Pooling equilibria at output levels 3-5 and separating equilibria with M_{LS} selecting 5-7

⁴The data set we used amalgamates samples from experiments conducted at the University of Pittsburgh from 1993 to 1998 and at the Ohio State University from 2001 to 2005. We have tried to add a dummy regressor to account for the possible different in the student population. But its coefficient is not significant at the conventional level.

involve strategic behavior (limit pricing), as M s produce above full-information levels.

[Figures 1 and 2 approximately here]

Figure 1 reports the results of games with both pure-strategy pooling and separating equilibria under the meaningful context treatment. Figure 2 reports the results of games with no pure-strategy pooling equilibria under the meaningful context treatment.⁵ Results are reported for each cycle (12 plays) of the game. The histograms show M s' choices, with entry rates given right below each output level. The experimental data show a gradual adjustment process starting with M s at their myopia maxima (2 for M_{HS} , 4 for M_{LS}) followed by an attempt to pool at 4 and then, if no pooling equilibrium exists, a separating equilibrium starting to develop. That is, both pooling and separating equilibria only develop in response to a gradual, history dependent adjustment/learning process.

Given that outcomes for all pairings (M s' choice, E s' response, and M s' type) are reported after each play of the game, as well as the nature of the game, there is clear scope for peer group effects in the data, the explicit study of which provides the primary motivation for the present paper. Given the gradual development of strategic play observed, it seems reasonable that M s learn to limit price in one of three ways (i) a trial and error learning process in response to E s' responses to their own choices, (ii) through observing E s' responses to other M s' choices, and (iii) through observing like-type M s' choices and responses to same. All three are avenues through which M s are likely to develop beliefs about responses to their choices, which beliefs lie at the heart of the theory.

To relate the present paper to the econometric literature on peer group effects, we follow Manski (1993) in specifying three main sources for peer group effects: similar behavior of individuals belonging to a given reference group may be due to *endogenous effects*, wherein “the propensity of an individual to behave in some way varies with the behavior of the group”; *exogenous effects*, wherein “the propensity of an individual to behave in some way varies with the exogenous characteristics of the group”; and *correlated effects*, wherein “individuals in the same group tend to behave similarly because they have similar individual characteristics or face similar institutional environments”. In experimental settings, exogenous effects and correlated effects can be largely controlled for through recruiting procedures and the experimental design. However, endogenous effects are relatively hard to control for as subjects typically interact with each other, with these effects likely to be most pronounced in experiments where players receive feedback about all other players' choices and responses

⁵The results under the generic context treatment show a similar pattern.

to same. This motivates the general dynamic model with social interactions introduced in the next section.

3 General Dynamic Discrete Choice Models with Social Interactions and Simulation Estimation

In what follows we propose a general dynamic discrete choice panel data model with lagged social interactions

$$y_{it}^* = h(y_{i,t-1}^*, \dots, y_{i,-\infty}^*, Y_{n,t-1}, \dots, Y_{n,-\infty}, X_{nt}, \dots, X_{n,-\infty}, \xi_i) + v_{it}, \quad (1)$$

for $i = 1, \dots, n$, where Y_{nt} is the n -dimensional vector of dichotomous indicators y_{1t}, \dots, y_{nt} such that $y_{it} = 1$ if $y_{it}^* > 0$ and $y_{it} = 0$ otherwise, and X_{nt} is the $n \times k$ -dimensional matrix of strictly exogenous variables. The terms $Y_{n,t-1}, \dots, Y_{n,-\infty}$ in $h(\cdot)$ capture the (lagged) endogenous social interactions effect, so that the current decision of a subject could depend on the past choices of peers. On the other hand, $X_{nt}, \dots, X_{n,-\infty}$ are introduced to capture exogenous (or contextual) effects, so that the current decision of a subject may also depend on the exogenous characteristics of peers. The random individual components ξ_i are i.i.d. $N(0, 1)$ for all i ⁶ and the disturbances v_{it} are i.i.d. $N(0, 1)$ for all i and t . This process starts at $t = 1$, and the initial conditions on y_{it}^* , Y_{nt} and X_{nt} for $t \leq 0$ are fixed outside the model and are assumed to be zero.⁷ The specification of the dynamic model in Heckman (1981) does not incorporate (lagged) social interactions in that $y_{i,s-1}$ and x_{is} appear but not $Y_{n,s-1}$ and X_{ns} ($s \leq t$). Depending on the specification of the function $h_{it}(\cdot)$ in terms of lagged observed or latent dependent variables, the Heckman discrete dynamic model is known to be sufficiently flexible to accommodate a wide variety of dynamic structures such as Markov models, Polya models, renewal processes, latent Markov models, with rich specifications on disturbances. It allows for unobserved heterogeneity across the cross-sectional units and serial correlation for the remaining disturbances. The model with social interactions in (1) is generalized to incorporate additional dynamic effects due to peers' actions.⁸

⁶As ξ_i is an unobservable, the coefficient of ξ_i will absorb its standard deviation. So we may just assume the unit variance of ξ_i here.

⁷This assumption is reasonable for the experimental data in our empirical analysis, as we can observe the data generating process from the very beginning.

⁸Hajivassiliou (2009) has generalized the model specification in Mundlak (1978) and Chamberlain (1984) to a dynamic model where the random individual component ξ_i might be correlated with some time-varying and/or time-

Let $h_{it} = h(y_{i,t-1}^*, \dots, y_{i,-\infty}^*, Y_{n,t-1}, \dots, Y_{n,-\infty}, X_{nt}, \dots, X_{n,-\infty}, \xi_i)$ for simplicity and Φ be the standard normal distribution function with density ϕ . Define the integral limits L_{it} and U_{it} , where $L_{it} = -h_{it}$ and $U_{it} = \infty$ if $y_{it} = 1$, and $L_{it} = -\infty$ and $U_{it} = -h_{it}$ otherwise. The joint probability of Y_{nT}, \dots, Y_{n1} conditional on exogenous variables is

$$\begin{aligned} & P(Y_{nT}, \dots, Y_{n1} | X_{nT}, \dots, X_{n1}) \\ &= \prod_{i=1}^n \int \dots \int \Phi(d_{iT} h_{iT}) [\prod_{s=1}^{T-1} \Phi(d_{i,T-s} h_{i,T-s}) \phi_{[L_{i,T-s}, U_{i,T-s}]}(v_{i,T-s}) dv_{i,T-s}] \phi(\xi_i) d\xi_i, \end{aligned}$$

where $d_{it} = 2y_{it} - 1$ and $\phi_{[L_t, U_t]}$ is a truncated standard normal density function with support $[L_t, U_t]$.⁹ This likelihood suggests that the GHK simulator can be recursively applied to construct a simulated likelihood. First, generate n independent standard normal random variables ξ_i . With initial conditions given, the random variables v_{it} can be generated in the following steps. For each i , from $t = 1$ to $T - 1$, (i) generate u_{it} from the uniform $[0, 1]$ distribution, (ii) compute $v_{it} = -d_{it} \Phi^{-1} [u_{it} \Phi(d_{it} h_{it})]$, and then (iii) generate the latent dependent variable $y_{it}^* = h_{it} + v_{it}$. With m independent simulation runs, the corresponding simulated log likelihood function is given by

$$\mathcal{L} = \sum_{i=1}^n \ln \left\{ \frac{1}{m} \sum_{j=1}^m \prod_{t=1}^T \Phi(d_{it} h_{it}^{(j)}) \right\}, \quad (2)$$

where $h_{it}^{(j)} = h_{it}(y_{i,t-1}^{*(j)}, \dots, y_{i1}^{*(j)}, Y_{n,t-1}, \dots, Y_{n1}, X_{nt}, \dots, X_{n1}, \xi_i^{(j)})$, and the superscript (j) denotes an independent simulation run. Thus, the simulation of the likelihood for the model (1), is similar to one of the conventional dynamic panel models in Lee (1997).

Asymptotic properties of the simulated maximum likelihood estimator (SMLE) for cross-sectional or short time series panel data have been studied in Hajivassiliou and McFadden (1990), Lee (1992; 1995) and Gourieroux and Monfort (1993), among others. The SMLE can be asymptotically efficient when m increases at a rate faster than $n^{1/2}$. However, as the likelihood function is nonlinear, the SMLE might have an asymptotic bias if the number of random draws to construct the likelihood simulator does not increase fast enough relative to the sample size. When m increases at a rate of $n^{1/2}$, as shown in Lee (1995), an asymptotic bias exists in the asymptotic distribution. The asymptotic bias will dominate the variance when m increases at a rate slower than $n^{1/2}$. Lee (1995)

invariant regressors. In the empirical model of this paper, all the explanatory variables are predetermined. So we assume that the individual component ξ_i can be normalized to have zero mean.

⁹Note that the dimension of integration does not depend on the number of peers because we assume that the current decision of a subject is only affected by the observed choices and characteristics of the peers in earlier periods.

has suggested a simple bias-correction procedure to remove the leading bias term due to simulation. The asymptotic efficiency of the bias-adjusted estimator requires only that m grows at a rate faster than $n^{1/4}$.¹⁰

4 Empirical Models and Estimation Results

In this section, we apply the dynamic discrete choice setting to investigate the evolution of subjects' behavior in the experiment. We estimate the model on two different samples: One from the experimental sessions with E_{HS} and the other from the sessions with E_{LS} . With E_{HS} , we model the learning process of M_{HS} , treating choices of 3-5 by M_{HS} as limit pricing as they learn to imitate M_{LS} and deter entry. For games with E_{LS} , we focus on strategic play of M_{LS} as they attempt to distinguish themselves from M_{HS} through choosing output levels 5-7, thereby deterring entry. Note that using stricter definitions of limit pricing such as choices of 6 and 7 by M_{LS} in games with E_{LS} or 4 and 5 for M_{HS} in games with E_{HS} yield similar results to those reported here.

As the adjustment (learning) process for the two samples are modeled analogously, we only detail the model specification for games with E_{HS} .

4.1 A Markov Model with Lagged Social Interactions

The unobservable incentives for M_{HS} to limit price are characterized by a Markov dynamic discrete choice model with lagged social interactions. The Markov model assumes a subject's current decision depends on his/her *last* decision and the feedback information from the *previous* round of the game. We will relax this restrictive assumption and consider a more general dynamic process in the next subsection.

By experimental design, a subject is randomly assigned turns as an M_H in different plays of the game within an experimental session. At the same time, a subject can observe peers' choices and E_S ' responses from all previous rounds of the game. As such, we distinguish between a *decision period* in which a subject plays as M_H with the opportunity to limit price and a (consecutive) *calendar period*. For a subject i , let T_i be the total number of decision periods in which he/she has played as an M_H . Corresponding to each decision period τ ($\tau = 1, \dots, T_i$), there is a calendar period. Let $t_i(\tau)$ be the

¹⁰We also conducted some Monte Carlo simulation experiments for the finite sample performance of the simulated estimator. An appendix containing the derivation of the likelihood function and simulation results is available upon request.

calendar period when subject i plays as M_H . The Markov model for subject i is specified as

$$y_{it_i(\tau)}^* = \alpha + \beta x_{i,t_i(\tau)-1} + \lambda_1 y_{i,t_i(\tau-1)} + \lambda_2 z_{i,t_i(\tau)-1} + \gamma \ln \tau + \sigma \xi_i + \varepsilon_{it_i(\tau)}, \quad (3)$$

for $\tau = 1, \dots, T_i$. We assume that $\varepsilon_{it_i(\tau)} = \rho \varepsilon_{i,t_i(\tau-1)} + v_{it_i(\tau)}$, and ξ_i and $v_{it_i(\tau)}$ are i.i.d. $N(0, 1)$. As the dynamic process starts at the first sampling period in the experiment, the initial conditions on all variables for $t \leq 0$ are zero. Other variables in (3) are explained below.

If the latent dependent variable $y_{it_i(\tau)}^* > 0$, subject i limit prices in their τ th turn as an M_H , and the corresponding dependent variable $y_{it_i(\tau)}$ is 1; $y_{it_i(\tau)}$ is 0 otherwise. Explanatory variables are on the right hand side of (3). α is a constant. $x_{i,t_i(\tau)-1}$ is the perceived entry rate differential between choices 1-2 and choices 3-5.¹¹ Specifically, let $d_{is}^L(IN)$ (respectively, $d_{is}^O(IN)$) be a dummy variable indicating that subject i chooses 3, 4 or 5 (respectively, 1 or 2) and observes the response IN in calendar period s . Let $d_{-is}^L(IN)$ (respectively, $d_{-is}^O(IN)$) be the number of M s other than i who, in calendar period s , choose 3, 4 or 5 (respectively, 1 or 2) and observe the response IN . Define $d_{is}^L(OUT)$, $d_{is}^O(OUT)$, $d_{-is}^L(OUT)$, and $d_{-is}^O(OUT)$ in an analogous manner, where OUT indicates potential E s staying out. Denote the weight a player puts on entry in response to other M s' choices by ω . The perceived entry rate differential is given by

$$x_{i,t_i(\tau)-1} = \frac{d_{i,t_i(\tau)-1}^O(IN) + \omega d_{-i,t_i(\tau)-1}^O(IN)}{d_{i,t_i(\tau)-1}^O + \omega d_{-i,t_i(\tau)-1}^O} - \frac{d_{i,t_i(\tau)-1}^L(IN) + \omega d_{-i,t_i(\tau)-1}^L(IN)}{d_{i,t_i(\tau)-1}^L + \omega d_{-i,t_i(\tau)-1}^L},$$

where $d_{i,t_i(\tau)-1}^j = d_{i,t_i(\tau)-1}^j(IN) + d_{i,t_i(\tau)-1}^j(OUT)$ for $j = L, O$.¹² This term serves as a proxy for the unobservable beliefs of M s regarding E s' responses to different choices. The time-lagged observed dependent variable $y_{i,t_i(\tau-1)}$ is introduced to measure the true state dependence of the dynamic process. The regressor $z_{i,t_i(\tau)-1}$, which is a function of $Y_{n,t_i(\tau)-1}$, denotes the proportion of like-type M s (other than i) limit pricing in the preceding calendar period. It captures the potential peer group effect resulting from observing other like-type M s' previous choices. The coefficient on $\ln \tau$, where τ is the number of decision periods that the subject i has played as M_H (including the current decision period), collects all other experience effects within an experimental session. A random individual component ξ_i is introduced to control for unobserved heterogeneity across

¹¹The entry rates are calculated conditional on the output level selected, not on the type of M , since E s can not observe M s' type when deciding to play IN or OUT.

¹²We assume that $(d_{i,t_i(\tau)-1}^j(IN) + \omega d_{-i,t_i(\tau)-1}^j(IN)) / (d_{i,t_i(\tau)-1}^j + \omega d_{-i,t_i(\tau)-1}^j) = 0.5$, in the case that $d_{i,t_i(\tau)-1}^j + \omega d_{-i,t_i(\tau)-1}^j = 0$ ($j = L, O$).

players. The remaining disturbances are assumed to follow an AR(1) process, as we find in a Monte Carlo study (available upon request) that flexible error specifications of this sort are favorable to identifying potential peer group effects.

We model the adjustment process of M_{LS} 's choices in experimental sessions with E_{LS} in an analogous manner, with $x_{i,t_i(\tau)-1}$ being the perceived entry rate differential between choices 3-4 and 5-7.

By a quasi-difference transformation of (3), i.e. by substituting $\varepsilon_{it_i(\tau)} = \rho[y_{it_i(\tau-1)}^* - (\alpha + x_{i,t_i(\tau-1)-1}\beta + \lambda_1 y_{i,t_i(\tau-2)} + \lambda_2 z_{i,t_i(\tau-1)-1} + \gamma \ln(\tau - 1) + \sigma \xi_i)] + v_{it_i(\tau)}$ in (3), it is easy to see that the empirical Markov model conforms to the general dynamic model (1). For subject i the likelihood function involves $(T_i - 1)$ -dimension integrals which are analytically intractable and numerically hard to evaluate. We circumvent this difficulty in the maximum likelihood estimation by a simulation method based on the unbiased GHK simulator. Table 5 reports the SMLEs for the Markov model based on simulations with 100 random draws using data from experimental sessions with E_{HS} and E_{LS} respectively.¹³

The positive and statistically significant SMLEs of λ_1 in all cases show that a subject's current choice depends heavily on his/her choice in the previous decision period, so that once strategic play is initiated, it substantially increases the likelihood of it continuing in later periods. This provides clear evidence of learning. Interaction terms (with the dummy variable NX representing sessions with inexperienced subjects) are introduced to account for the differences between experienced and inexperienced subject sessions.¹⁴ The negative and significant coefficient estimate for the interaction term of lagged choice and NX in games with E_{LS} indicates that inexperienced subjects were much less likely to play strategically having once chosen to do so than their experienced selves. We interpret this as indicating that inexperienced M s were less confident regarding their choice of strategic play in games with E_{LS} , than in games with E_{HS} , due to the more challenging nature of the game.

Peer group effects are modeled in two distinct ways in the Markov model: (i) through the entry rate differential in response to own choices and to other M s' choices and (ii) through $z_{i,t_i(\tau)-1}$ which captures the effect of observing like-type M s' choices. Exploring which of these two factors will impact an M 's strategic choice and their relative importance provides much of the focus for the

¹³We have tried to add more interaction terms, or remove some regressors or interaction terms with insignificant coefficients. The estimation results are robust to these different specifications.

¹⁴The minus two times log likelihood ratios for testing jointly the significance of those interaction terms in the Markov model are, respectively, 8.78 for games with E_{HS} , and 8.62 for games with E_{LS} . Neither is significant at the 5 percent level.

analysis that follows.

For games with E_{HS} , none of the peer group variables by itself has a statistically significant impact on M_{HS} 's likelihood of limit pricing. Further, none of the peer group coefficients taken jointly (β , ω , and/or λ_2) are statistically significant at conventional levels. However, as will be reported in the next subsection, peer group effects are identified in the Polya model.

In games with E_{LS} , the SMLEs of λ_2 and β are positive and statistically significant, indicating that both factors have positive impact on an M_L 's decision to limit price. For the specification without interaction terms, the average marginal impact of other M_{LS} 's choices (λ_2) on the probability of limit pricing is 0.043.¹⁵ The average marginal impact of the entry rate differential on the probability of limit pricing is smaller, 0.033. However, the estimated ω is significantly less than 1, indicating M_{LS} place primary weight on what happens to themselves, with very limited weight placed on what happens to others when they attempt to separate; e.g., in the specification without interaction terms, M_{LS} place about 67 times more weight on own entry than on entry to another M making the same choice (see Simonsohn et al., in press, for an experiment reporting similar severe overweighting regarding own outcomes versus others in a prisoner's dilemma experiment when, for all practical purposes, they should be weighted equally). Further, the small weight placed on entry in response to other M s choices cannot be attributed to the presence of $z_{i,t_i(\tau)-1}$, as dropping it has little impact on the coefficient value for ω .

The proportion of M_{HS} attempting to pool by choosing 3 and 4 in the previous round is introduced as an additional explanatory variable in games with E_{LS} because increased frequencies on this score generates increased entry rates on 3 and 4, which in turn motivates M_{LS} to separate. Although positive in sign, its coefficient fails to achieve statistical significance in the Markov specification.

The positive and significant estimates of the coefficient on $\ln \tau$ in games with E_{HS} picks up other experience effects that fail to be captured in the model. There is a strong interaction effect with NX indicating that these factors are largely confined to experimental sessions with inexperienced

¹⁵For the general model (1), $E(y_{it}|(y_{is}^*, Y_{ns}, X_{ns}, s = 1, \dots, t-1), X_{nt}, \xi_i) = \Phi(h_{it})$. The average marginal effect over time and individual of, say x_{it} (which is assumed continuous), on the transition probability $P(y_{it} = 1|Y_{ns}, X_{ns}, s = 1, \dots, t-1), X_{nt})$, is given by

$$\frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T \int \dots \int \phi(h_{it}) (\partial h_{it} / \partial x_{it}) f(y_{i1}^*, \dots, y_{i,t-1}^*, \xi_i | Y_{ns}, X_{ns}, s = 1, \dots, t-1) dy_{i1}^* \dots dy_{i,t-1}^* d\xi_i$$

The multiple integrals here can be approximated by simulations. We simulate $h_{it}^{(j)}$ following the same procedure as in (2). With m independent simulation runs, the corresponding (sample average) simulated marginal effects is $\sum_{j=1}^m \sum_{i=1}^n \sum_{t=1}^T \phi(h_{it}^{(j)}) (\partial h_{it} / \partial x_{it}) / mnT$. Results reported in this paper are based on a simulator generated from 1000 random draws.

subjects. These unexplained learning effects are one factor motivating the development of a more general empirical model in the next subsection.

The estimated coefficients of the dummies for experienced players are positive and statistically significant in games with E_{HS} and with E_{LS} , indicating that in both cases experienced subjects start out with much higher levels of strategic play than their inexperienced selves. The negative and significant estimates for the constants (α) indicate that it takes time for strategic play to develop in all cases. The larger absolute value for α in games with E_{LS} indicates that strategic play is much slower to emerge in this case, a fact that is quite apparent in the raw data.

Though the overall correlation in the disturbances captured by $\sigma\xi_i + \varepsilon_{it_i(\tau)}$ is positive, the estimated AR(1) parameter ρ is negative. In a Monte Carlo study (available upon request), we found that the estimator of ρ can be strongly downward biased when a Polya process is misspecified as a Markov process. Hence, the negative $\hat{\rho}$ could be a downward biased estimate due to model misspecification. In the next subsection, we generalize the Markov model to a more general dynamic process.

4.2 A Polya Model with Lagged Social Interactions

As subjects have access to all previous outcomes in an experimental session, the entire history of past plays could be relevant to the current choices. In this subsection, we model the influence of all past plays on a subject's current decision employing a Polya process with lagged social interactions. Similarly to the Markov model, we assume that the unobservable incentives to limit price can be characterized by

$$y_{it_i(\tau)}^* = \alpha + \beta \bar{x}_{i,t_i(\tau)-1} + \lambda_1 \sum_{s=1}^{\tau} \delta_1^{s-1} y_{i,t_i(\tau-s)} + \lambda_2 \frac{\sum_{s=1}^{t_i(\tau)} \delta_2^{s-1} z_{i,t_i(\tau)-s}}{\sum_{s=1}^{t_i(\tau)} \delta_2^{s-1}} + \gamma \ln \tau + \sigma \xi_i + \varepsilon_{it_i(\tau)}, \quad (4)$$

and $\varepsilon_{it_i(\tau)} = \rho \varepsilon_{i,t_i(\tau-1)} + v_{it_i(\tau)}$, where $\xi_i, v_{it_i(\tau)}$ are i.i.d. $N(0,1)$. The the initial conditions on all variables for $t \leq 0$ set to be zero. Most variables in (4) are defined as in the Markov model (3), except the cumulative entry rate differential $\bar{x}_{i,t_i(\tau)-1}$. Let $c_{i,t_i(\tau)-1}^j(R) = \sum_{s=1}^{t_i(\tau)-1} d_{is}^j(R)$ and $c_{-i,t_i(\tau)-1}^j(R) = \sum_{s=1}^{t_i(\tau)-1} d_{-is}^j(R)$ for $j = L, O$ and $R = IN, OUT$, with $d_{is}^j(R)$ given as before. The perceived cumulative entry rate differential between ‘‘myopic’’ choices and strategic choices is

given by

$$\bar{x}_{i,t_i(\tau)-1} = \frac{c_{i,t_i(\tau)-1}^O(IN) + \omega c_{-i,t_i(\tau)-1}^O(IN)}{c_{i,t_i(\tau)-1}^O + \omega c_{-i,t_i(\tau)-1}^O} - \frac{c_{i,t_i(\tau)-1}^L(IN) + \omega c_{-i,t_i(\tau)-1}^L(IN)}{c_{i,t_i(\tau)-1}^L + \omega c_{-i,t_i(\tau)-1}^L},$$

where $c_{i,t_i(\tau)-1}^j = c_{i,t_i(\tau)-1}^j(IN) + c_{i,t_i(\tau)-1}^j(OUT)$ for $j = L, O$, with ω indicating the weight a player puts on the entries on other M s relative to his/her own. Similar to $x_{i,t_i(\tau)-1}$ in the Markov model (3), $\bar{x}_{i,t_i(\tau)-1}$ here represents the payoff incentive for M to limit price.¹⁶ The depreciation factors δ_1 and δ_2 measure the influence of past plays on current choices. The coefficient on the weighted average $\sum_{s=1}^{t_i(\tau)} \delta_2^{s-1} z_{i,t_i(\tau)-s} / \sum_{s=1}^{t_i(\tau)} \delta_2^{s-1}$ captures the cumulative effect of like-type M s' past choices. Thus, in the Polya model, a subject's current decision is assumed to be influenced by the (weighted) average of their peers' choices over the entire history of play. As in the Markov model, for games with E_{LS} , we introduce the proportion of M_{HS} attempting to pool by choosing output levels 3 and 4 as an additional explanatory variable. Unlike the Markov model where this value is calculated based on M_{HS} ' choices in the previous round only, we calculate its cumulative counterpart based on all past rounds in the Polya model. Based on the GHK simulator with 100 simulation draws, the SMLEs of the Polya model using data from the experimental sessions with E_{HS} and E_{LS} are reported in Tables 6 and 7 respectively.

In games with E_{HS} and E_{LS} , the positive and significant estimates of λ_1 on own past choices imply that previous strategic play substantially increases the likelihood of current strategic play. The coefficient on the interaction term between own past choices and NX (the dummy for sessions with inexperienced subjects) is not statistically significant in games with E_{HS} but is negative and statistically significant in games with E_{LS} .¹⁷ This is similar to the Markov model, indicating that in more difficult games (with E_{LS}), inexperienced subjects are more likely to revert back to non-strategic play than their more experienced selves.

In games with E_{HS} , the effect of peers' past choices in model (iii) of Table 6 is positive and statistically significant for inexperienced subjects (with a t -ratio of 2.95), with no significant additional learning from peers' choice for experienced players. In contrast, the coefficient estimate for

¹⁶We have tried to introduce a time depreciation factor to the specification of the entry rate differential. However, the estimates of ω and that depreciation factor both have huge standard errors. As an alternative, we have tried to assign a predetermined value to the depreciation factor (like 0.2, 0.6, and 1). The estimation results are quite robust to the predetermined value of the depreciation factor.

¹⁷The interaction terms in the Polya model are jointly significant at the 5 percent level with the minus two times log likelihood ratios being 8.86 for the games with E_{HS} , and significant at the 1 percent level with the minus two times log likelihood ratios being 22.06 for games with E_{LS} .

the cumulative entry rate differential is not statistically significant, which also makes the estimate of ω extremely imprecise.¹⁸

In games with E_{LS} , cumulative effects of other M_{LS} ' choices are positive and statistically significant, with even stronger effects for inexperienced subjects as indicated by the positive coefficient estimate for the interaction term between other M_{LS} ' choices and NX (with a t -ratio of 1.91). Thus, inexperienced M_{LS} are influenced more than experienced M_{LS} by their peers' choices. But unlike in games with E_{HS} , experienced subjects continue to be influenced by their peers, indicative of the fact that it takes longer for a separating than a pooling equilibrium to emerge. In addition, the coefficient on the cumulative entry rate differential is statistically significant in games with E_{LS} . The marginal effect of the entry rate differential is essentially the same as the impact of other M_{LS} choices for experienced subjects, but not so for inexperienced players. Rather, for inexperienced M_{LS} , what drives limit pricing is the frequency with which M_{HS} chose 3-4 as the coefficient value here is among the largest reported with a t -ratio of 2.94. The increased frequency of M_{HS} attempting to pool raises entry rates on 3-4, providing a clear message to M_{LS} that it's no longer profit maximizing to stick with 4, thereby driving them to choose higher output levels in order to distinguish themselves.

In games with E_{HS} and with E_{LS} , the SMLEs for both δ_1 and δ_2 are positive and statistically significant, indicating that a subject's current decisions are influenced by all past plays of the game, albeit with some depreciation as both δ s are less than one.¹⁹ Other experience effects as represented by $\ln \tau$ are not statistically significant at conventional levels in the Polya model.²⁰ Further, the sign of ρ in the AR(1) disturbances, which is negative in the Markov model, now becomes positive.

One element that has been left out of the analysis so far involves distinguishing between attempts at limit pricing that have successfully deferred entry as opposed to those that have not. To capture this we introduce two new variables into the regressions: an M 's own successful limit pricing and the percentage of successful limit pricing by peers. We view these new regressors as, essentially, interaction terms between M s' choices and E s' responses. The results are reported in the last columns of Tables 6 and 7. Introduction of these variables has essentially no effect on the estimation results for games with E_{HS} . For games with E_{LS} , the new regressors are jointly significant at the 1%

¹⁸Note that ω would not be identifiable if the coefficient of x_{it} were zero.

¹⁹Note, δ_1 and δ_2 are not directly comparable, as the depreciation factor δ_1 for own lagged choices is defined on the decision period whereas the depreciation factor δ_2 for like-type M s choices are defined on the (consecutive) calendar period. In addition δ_1 is associated with a weighted sum of past choices whereas δ_2 is associated with a weighted average, so that they are not comparable on that dimension either.

²⁰It seems plausible that the experience of being the type that can limit price is more important than that of being the other type. We have experimented with an additional regressor capturing the experience of being the other type of M . But the coefficient of that regressor failed to be significant at the conventional level.

level. Own success at limit pricing plays a statistically significant role in promoting limit pricing, and reduces the importance of own past attempts at limit pricing (λ_1). Not surprisingly, M_{LS} receive more encouragement to play strategically following success than failure. The effect of the cumulative entry rate differential diminishes, but the frequency of M_{HS} attempting to pool stays significantly positive in sessions with inexperienced subjects, with a t -ratio of 3.49. The successful limit pricing by peers plays no independent, statistically significant, role in promoting limit pricing, while peers attempts' at limit pricing maintains its statistical significance. This is consistent with the low estimated weight for entry on peers (ω) as M_{LS} tend to ignore E s responses to peers' limit pricing.

The first columns in Tables 6 and 7 consider the impact of neglecting peers' past choices. For games with E_{HS} , there is little if any effect on any of the estimates and only a small change in the log likelihood value. For games with E_{LS} , the SMLEs for the coefficient of the entry rate differential (β) and the weight of entries on peers (ω) are affected by dropping peers' past choices most, with upward biases. This is not too surprising since it is the entry rate differential between myopic and strategic choices that drive M_{LS} to limit price in the first place. In this context what the introduction of peers' past choices does is to better identify what information agents are focusing on that generates peer group effects in these games.

4.3 Model Selection and Cost of Learning

We prefer the Polya model to the Markov model for the following reasons. First, the Polya model is more consistent with a priori considerations that unless payoffs have changed, its unreasonable to think that subjects rely strictly on what happened in the previous period only. Second, the estimated AR(1) coefficient in the error term of the Markov model is negative while that of the Polya model is positive. This is consistent with our finding that in Monte Carlo experiments the estimate of the AR(1) coefficient in the error term will be downward biased when a Polya model is misspecified as a Markov process.²¹ Further, although the Polya model does not nest the Markov model because of the different specifications of some regressors, the well known Akaike information criterion (AIC) indicates the Polya model is better than the Markov model as the former has a smaller AIC value.²²

²¹ These simulation results are available on request.

²² The AIC is given as $AIC = -\frac{2}{n} \log L + \frac{2k}{n}$, where n is the sample size, k is the number of parameters and $\log L$ is log likelihood of a model. For the specification without interaction terms, in games with E_{HS} , the AIC of the Markov model is 0.9890 and the AIC of the Polya model is 0.9848; and in games with E_{LS} , the AIC of the Markov model is 0.7322 and the AIC of the Polya model is 0.7209.

Two questions not addressed to this point are (i) what is the financial cost to players resulting from the adaptive learning process as opposed to limit pricing, and (ii) what are the implied benefits of attending to the peer group effects associated with observing like type M s choices? To answer these questions, we conduct a limited simulation study. Given the observed entry rates in the raw data, we generate M s' choice using the estimated Markov and Polya models. We set the initial conditions of all variables in the dynamic process to the observed values in the first period of the raw data. For the Markov model, we use the estimates without interaction terms with the dummy variables in Table 5. For the Polya model, we use the estimates for model/specification (ii) in Tables 6 and 7.²³ To calculate M 's payoff, we assume that M_H chooses 4 when limit pricing and 2 otherwise in games with E_{HS} . In games with E_{LS} , we assume that M_L chooses 6 when limit pricing and 4 otherwise.²⁴ The simulation results based on 200 repetitions are reported in Tables 8 and 9. The simulated limit pricing rates based on the estimated Markov and Polya models are reported in column (b) of Table 8. The simulated limit pricing rates for both models are quite close to the actual rate in the raw data.²⁵

[Tables 8 and 9 approximately here]

Regarding (i), we look at the expected revenue loss from the adaptive learning (i.e. column (c) of Table 9) relative to the expected payoff of limit pricing (i.e. column (a) of Table 9).²⁶ This is minimal to begin with under both models mostly because it takes some time for E s' entry decisions to converge to the equilibrium prediction. So we focus on the expected revenue loss in later plays of an experimental session. Instead of following the Markov learning process, M_{HS} would have earned 13.8% (7.7%) more as inexperienced (experienced) subjects in the last 12 plays (cycle 3) of games with E_{HS} by always limit pricing, and M_{LS} would have earned 7.8% (6.3%) more as inexperienced (experienced) subjects in cycle 3 of games with E_{LS} by always limit pricing. On the other hand, instead of following the Polya learning process, M_{HS} would have earned 12.7% (6.2%) more as inexperienced (experienced) subjects in cycle 3 of games with E_{HS} by always limit pricing, and M_{LS}

²³The simulation results based on the estimates with interactions terms with dummy regressors are similar.

²⁴Calculations based on choice of 3 versus 2 in games with E_{HS} and those based on choice of 5 versus 4 in games with E_{LS} are similar to what are reported here.

²⁵For both models, the biggest deviation of the simulated rates from the actual rate occurs in the last 12 plays (cycle 3) of the experienced-subject sessions. This is mainly because most experienced-subject sessions only last 24 plays (2 cycles). There are only a small number of observations for cycle 3 in those sessions.

²⁶One limitation on the following counterfactual study is E s' choices are based on the experimental data, whereas more limit pricing would be expected to impact E s' choices. However, the fact that E s do not play equilibrium in the data (e.g., always entering on 4 and always staying out on 6 in games with E_{LS}) reduces the losses resulting from the adaptive learning process versus strict limit pricing.

would have earned 7.4% (5.4%) more as inexperienced (experienced) subjects in cycle 3 of games with E_{LS} by always limit pricing.²⁷ M s are slightly better off by following the Polya process than following the Markov process.

With respect to the question (ii) above, we look at the average difference in limit pricing rate and payoff of simulated learning process with and without like-type M peer group effects (i.e. columns (b) and (c) of Table 8 and columns (c) and (d) of Table 9). In cycle 3 of games with E_{HS} , the average difference in the rate of limit pricing is less than 2% and the average difference in payoffs is less than 1% for both inexperienced and experienced subjects under the Markov learning process. The average difference in limit pricing is 18.8% (13.4%) and the average difference in payoffs is 3.6% (3.7%) for inexperienced (experienced) subjects under the Polya learning process. In cycle 3 of games with E_{LS} , for inexperienced subjects, the average difference in limit pricing is 17.4% (42.1%) under the Markov (Polya) learning process, but the average difference in payoffs is minimal (less than 1%) in both cases. For experienced players, the average difference in limit pricing is 23.2% (46.5%) and the average difference in payoffs is 3.3% (6.0%) under the Markov (Polya) learning process. The reason limit pricing is not that rewarding in payoff as it should be is that the observed entry rate in the raw data converges to the equilibrium prediction quite slowly, especially in game with E_{LS} . For example, in cycle 3 of an inexperienced-subject (experienced-subject) session with E_{LS} , the entry rate on choice 4 (i.e., not limit pricing) is 59.3% (68.2%) instead of 100% predicted by the Nash equilibrium.

5 Summary and Conclusions

We have investigated learning and peer group effects in laboratory experiments based on Milgrom and Roberts' (1982) entry limit pricing game. We employed both Markov and Polya models with lagged social interactions to characterize the development of strategic play over time. Compared to the Markov model, the Polya model has a more natural justification and provides a slightly better fit to the data in terms of the AIC criteria. Both models tell the same story with respect to peer group effects in that entry rates play a significant role in motivating strategic play as does strategic

²⁷Translated into the dollar payoffs employed in the Ohio State sessions, for the Markov model these amount to losses per play of the game of 11.2 cents (7.2 cents) for inexperienced (experienced) subjects in games with E_{HS} and 9.3 cents (8.3 cents) per play of the game for inexperienced (experienced) players in games with E_{LS} . For the Polya model these amount to losses of 10.6 cents (5.9 cents) for inexperienced (experienced) subjects in games with E_{HS} and 8.9 cents (7.3 cents) per play of the game for inexperienced (experienced) players in games with E_{LS} .

play of like-type monopolists.

From a methodological point of view the dynamic panel data model allows us to study the adjustment process in substantially more detail than past studies that have typically employed panel data models. The latter allows one to determine correlations between the aggregate frequency of strategic play and subjects' experience but reveals very little about the detailed nature of the learning process. However, there is a trade-off between the more detailed modeling employed in the present study and panel data models with time dummies in that the present analysis typically requires substantially more data than the latter. In the present case we have employed data from a large number of experimental sessions collected from several different papers. To get some idea of the data session requirements to estimate the present model, we have played with estimating the model for games with E_L s under several natural ways of breaking up the data to get some idea of the sensitivity of our results to the number of experimental sessions available to work with. Results are sensitive to including data from inexperienced versus experienced sessions: it is easier to identify the role of like-type monopolists' choices with more inexperienced subject sessions, and easier to identify the role of entrants' responses with more experienced subject sessions. What we can say is that sometimes we can distinguish between these two effects with as few as 7-10 experimental sessions, but that in general it is far better to have more sessions. As such it may be impractical to employ dynamic discrete choices learning models with social interactions within any given experimental study as the data requirements are too intense, but are rather more appropriately employed in a meta analysis of several different experimental studies.

The main substantive finding of the present study relates to peer group effects in learning. The analysis shows that subjects respond to several different types of peer group effects: (i) opponents' (entrants') behavior in the game (both in terms of responses to a player's own choices and to other players' choices) and (ii) peers' (like-type monopolists') past decisions. One would expect agents to respond to opponents' behavior in strategic interactions. From a theoretical perspective, one would expect that agents would (or at least should) place more weight on the population sample data than on one's own experience, as the former is far more informative. Yet this is not what happens with respect to direct observation of opponents choices, where the weight on the entry rate differential for the population sample (ω) is small in size and not significantly different from zero in both models. This is not entirely unanticipated from a behavioral perspective. In a recent paper, Simonsohn et al. (2008) look at repeated game experiments in which groups of players are randomly

rematched in every round and receive feedback about the actions and outcomes of all players. They find that, in a repeated prisoner's dilemma game, players are substantially more likely to cooperate when another player has cooperated with them in the last period than to the overall percentage of cooperators in last period. Further, the attention paid to percentage of other cooperators is substantially more mixed than to their own experience as the former does not rise to statistical significance at conventional levels, while the latter does.²⁸

One question that has not been addressed so far is if the peer group effect of observing what like-type monopolists are doing is a pure imitation effect, or an alternative/additional measure players use to determine the likelihood of success at limit pricing. Although we are not able to answer this question on the basis of the data reported here, a subsequent experiment using two person teams as decision makers suggests that something more than pure imitation is at work. In this experiment team dialogues were analyzed and coded into a number of different categories to get some insight into players' decision processes. First, consistent with the peer group effects identified here, by far the highest single coded category related to the information feedback players got regarding other *Ms* choices and responses to same (Cooper and Kagel, 2005). Another frequently coded category captured how teammates justified strategic play to each other. The most frequently coded subcategory here consisted of *Ms* explicitly reasoning from the point of view of *Es* potential responses to their choices (reasoning from a game theoretic point of view).²⁹ Having a conversation of this sort proved to be a very good predictor of whether or not, having played strategically once, a team continued to play strategically. This suggests that while looking at peer *Ms* choices, and *Es* responses to same, may motivate strategic play initially, it takes some understanding of the underlying strategic interactions, to generate continued strategic play.

References

Borsch-Supan, A. and Hajivassiliou, V. (1993). Smooth unbiased multivariate probability simulators for maximum likelihood estimation of limited dependent variable models, *Journal of Econometrics* **58**: 347–368.

²⁸In this respect also see Garvin and Kagel (1994) who show that the main factor teaching bidders to overcome the winner's curse in common value auction experiments is the "hot stove" effect of actually losing money as opposed to less painful observational learning.

²⁹For example: "if we enter 6 when we are an A2 (an M_L) everyone will know that we are an A2 and will guess accordingly, giving us a higher average....The 6 pays less but is very clear what we are."

- Butler, J. S. and Moffitt, R. (1982). A computationally efficient quadrature procedure for the one-factor multinomial probit model, *Econometrica* **50**: 761–764.
- Chamberlain, G. (1984). Panel data, in Z. Griliches and M. D. Intriligator (eds), *Handbook of Econometrics*, Vol. 2, Elsevier, Amsterdam.
- Cho, I. and Kreps, D. (1987). Signaling games and stable equilibria, *Quarterly Journal of Economics* **102**: 179–221.
- Cooper, D. J., Garvin, S. and Kagel, J. H. (1997a). Adaptive learning vs. equilibrium refinements in an entry limit pricing game, *The Economic Journal* **107**: 553–575.
- Cooper, D. J., Garvin, S. and Kagel, J. H. (1997b). Signaling and adaptive learning in an entry limit pricing game, *RAND Journal of Economics* **28**: 662–683.
- Cooper, D. J. and Kagel, J. H. (2003a). The impact of meaningful context on strategic play in signaling games, *Journal of Economic Behavior and Organization* **50**: 311–337.
- Cooper, D. J. and Kagel, J. H. (2003b). Lessons learned: generalizing learning across games, *The American Economic Review* **93**: 202–207.
- Cooper, D. J. and Kagel, J. H. (2004). Learning and transfer in signaling games. Working Paper.
- Cooper, D. J. and Kagel, J. H. (2005). Are two heads better than one? team versus individual play in signaling games, *The American Economic Review* **95**: 477–509.
- Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*, MIT Press, Cambridge.
- Garvin, S. and Kagel, J. H. (1994). Learning in common value auctions: some initial observations, *Journal of Economic Behavior and Organization* **25**: 351–372.
- Geweke, J. (1991). Efficient simulation from the multivariate normal and student-t distributions subject to linear constraints, *Computer science and statistics: Proceedings of the Twenty-Third Symposium on the Inference*, American Statistical Association, Alexandria, VA, pp. 571–578.
- Geweke, J. and Keane, M. (2001). Computationally intensive methods for integration in econometrics, in J. J. Heckman and E. Leamer (eds), *Handbook of Econometrics*, Vol. 5, Elsevier, Amsterdam.

- Gourieroux, C. and Monfort, A. (1993). Simulation-based inference: a survey with special reference to panel data models, *Journal of Econometrics* **59**: 5–33.
- Hajivassiliou, V. (2009). Estimation and specification testing of panel data models with non-ignorable persistent heterogeneity. LSE working paper.
- Hajivassiliou, V. and McFadden, D. (1990). The method of simulated scores for the estimation of ldv models: with an application to the problem of external debt crises. Cowles foundation discussion paper No. 967. Yale University, New Haven, CT.
- Heckman, J. J. (1981). Statistical models for discrete panel data, *in* C. F. Manski and D. McFadden (eds), *Structural Analysis of Discrete Data with Econometric Applications*, MIT Press, Cambridge.
- Hsiao, C. (2003). *Analysis of Panel Data*, Cambridge University Press, Cambridge.
- Keane, M. P. (1994). A computationally practical simulation estimator for panel data, *Econometrica* **62**: 95–116.
- Lee, L. F. (1992). On efficiency of methods of simulated moments and maximum simulated likelihood estimation of discrete response models, *Econometric Theory* **8**: 518–552.
- Lee, L. F. (1995). Asymptotic bias in maximum simulated likelihood estimation of discrete choice models, *Econometric Theory* **11**: 437–483.
- Lee, L. F. (1997). Simulated maximum likelihood estimation of dynamic discrete choice statistical models: some monte carlo results, *Journal of Econometrics* **82**: 1–35.
- Lee, L. F. (2000). A numerically stable quadrature procedure for the one-factor random-component discrete choice model, *Journal of Econometrics* **95**: 117–129.
- Lee, L. F. (2001). Interpolation, quadrature, and stochastic integration, *Econometric Theory* **17**: 933–961.
- Manski, C. F. (1993). Identification of endogenous social effects: the reflection problem, *The Review of Economic Studies* **60**: 531–542.
- Milgrom, P. and Roberts, J. (1982). Limit pricing and entry under incomplete information: an equilibrium analysis, *Econometrica* **50**: 443–459.

- Miller, R. M. and Plott, C. R. (1985). Product quality signaling in experimental markets, *Econometrica* **53**: 837–872.
- Mundlak, Y. (1978). On the pooling of time series and cross section data, *Econometrica* **46**: 69–85.
- Simonsohn, U., Karlsson, N., Loewenstein, G. and Ariely, D. (2008). The tree of experience in the forest of information: overweighing experienced relative to observed information, *Games and Economic Behavior* **62**: 263–286.

Table 1: A Monopolist's Payoffs

M_H (High Cost M)		M_L (Low Cost M)			
Your Choice	X (In)	Y (Out)	X (In)	Y (Out)	Your Choice
1	150	426	250	542	1
2	168	444	276	568	2
3	150	426	330	606	3
4	132	408	352	628	4
5	56	182	334	610	5
6	-188	-38	316	592	6
7	-292	-126	213	486	7

Source: Cooper et al. (1997b).

Table 2: A High Cost Entrant's Payoffs

M Player's Type			
Your Action Choice	M_H	M_L	Expected Value ^a
	(High Cost M)	(Low Cost M)	
X (In)	300	74	187
Y (Out)	250	250	250

^a Based on prior distribution (50% M_H , 50% M_L) of M types.

Source: Cooper et al. (1997b).

Table 3: A Low Cost Entrant's Payoffs

<i>M</i> Player's Type			
Your Action	<i>M_H</i>	<i>M_L</i>	Expected Value ^a
	(High Cost <i>M</i>)	(Low Cost <i>M</i>)	
Choice	Your Payoff	Your Payoff	
X (In)	500	200	350
Y (Out)	250	250	250

^a Based on prior distribution (50% M_H , 50% M_L) of M types.

Source: Cooper et al. (1997b).

Table 4: Experimental Treatments

Payoff Tables	Number of Sessions		Location	Dates	BNE Prediction
	<i>GC</i> ^a	<i>MC</i>			
1 & 2 (E_H s)	2 (30) / 0 (0) ^b	5 (72) / 4 (41)	Pittsburgh	03/93-03/98	Pure-strategy pooling & separating equilibria exist
	5 (70) / 4 (50)	4 (60) / 3 (44)	Columbus	04/01-03/02	
1 & 3 (E_L s)	7 (108) / 4 (62)	8 (110) / 5 (66)	Pittsburgh	03/93-03/98	No pure-strategy pooling equilibria exist
	8 (128) / 7 (100)	4 (62) / 4 (54)	Columbus	03/01-03/05	

^aGC: generic context; MC: meaningful context.

^bnumber of inexperienced-subject sessions / number of experienced-subject sessions; number of subjects in parentheses.

Table 5: SMLEs for the Markov Model

	High-Cost Type Entrants*		Low-Cost Type Entrants**	
		w/ interactions		w/ interactions
entry rate differential (β)	0.105 (.084)	0.184 (.182)	0.335 (.072) ^c	0.389 (.106) ^c
entry rate differential $\times NX$	-	-0.154 (.195)	-	-0.070 (.123)
weight of entries on peers (ω)	0.001 (.261)	0.525 (2.84)	0.015 (.076)	0.018 (.079)
own past choices (λ_1)	0.829 (.181) ^c	0.811 (.238) ^c	1.768 (.101) ^c	2.033 (.130) ^b
own past choices $\times NX$	-	0.137 (.209)	-	-0.397 (.141) ^c
peers' past choices (λ_2)	0.050 (.127)	-0.071 (.285)	0.434 (.118) ^c	0.365 (.163) ^b
peers' past choices $\times NX$	-	0.113 (.317)	-	0.121 (.236)
% of M_H choosing 3,4	-	-	0.038 (.099)	-0.105 (.183)
% of M_H choosing 3,4 $\times NX$	-	-	-	0.226 (.226)
experience in a session (γ)	0.222 (.080) ^c	-0.119 (.211)	0.052 (.048)	-0.014 (.077)
experience in a session $\times NX$	-	0.387 (.206)	-	0.080 (.087)
constant (α)	-0.554 (.154) ^c	-0.627 (.153) ^c	-1.705 (.097) ^c	-1.731 (.106) ^c
random effects (σ)	1.100 (.100) ^c	1.061 (.096) ^c	0.763 (.061) ^c	0.752 (.062) ^c
serial correlation (ρ)	-0.139 (.099)	-0.189 (.097) ^a	-0.260 (.051) ^c	-0.250 (.055) ^c
Dummies:				
sessions w/ experienced players	0.674 (.132) ^c	1.015 (.236) ^c	0.790 (.078) ^c	0.883 (.145) ^c
sessions w/ meaningful context	-0.089 (.160)	-0.080 (.158)	-0.082 (.094)	-0.075 (.094)
Log Likelihood	-1070.53	-1066.14	-1663.94	-1659.63

NX is a dummy variable for experimental sessions employing subjects with no experience of the same or related games

* 2185 observations on M_H ; ** 4575 observations on M_L

^a significantly different from 0 at the 10 percent level

^b significantly different from 0 at the 5 percent level

^c significantly different from 0 at the 1 percent level

standard errors in parentheses

Table 6: SMLEs for the Polya Model (Experiments with High-Cost Type Entrants)

	model (i)	model (ii)	model (iii)	model (iv)
cumulative entry rate diff. (β)	0.074 (.135)	0.046 (.135)	0.241 (.212)	0.173 (.180)
cumulative entry rate diff. $\times NX$	-	-	-0.227 (.263)	-
weight of entries on peers (ω)	0.528 (5.42)	0.556 (9.20)	0.652 (5.94)	0.947 (5.06)
own past choices (λ_1)	0.530 (.164) ^c	0.554 (.163) ^c	0.470 (.182) ^c	0.498 (.179) ^c
own past choices $\times NX$	-	-	0.152 (.142)	-
own successful limit pricing	-	-	-	0.083 (.107)
depreciation factor (δ_1)	0.593 (.150) ^c	0.585 (.147) ^c	0.512 (.170) ^c	0.581 (.149) ^c
peers' past choices (λ_2)	-	0.464 (.252) ^a	-0.387 (.519)	0.797 (.360) ^b
peers' past choices $\times NX$	-	-	1.290 (.616) ^b	-
peers' successful limit pricing	-	-	-	-0.565 (.445)
depreciation factor (δ_2)	-	0.929 (.321) ^c	0.990 (.194) ^c	0.895 (.271) ^c
experience in a session (γ)	0.103 (.109)	0.040 (.111)	0.043 (.110)	0.044 (.113)
constant (α)	-0.401 (.141) ^c	-0.534 (.156) ^c	-0.719 (.168) ^c	-0.544 (.157) ^c
random effects (σ)	0.963 (.126) ^c	0.941 (.124) ^c	0.961 (.122) ^c	0.938 (.126) ^c
serial correlation (ρ)	0.103 (.107)	0.084 (.106)	0.042 (.118)	0.088 (.107)
Dummies:				
sessions w/ experienced players	0.624 (.138) ^c	0.541 (.152) ^c	1.217 (.314) ^c	0.556 (.150) ^c
sessions w/ meaningful context	-0.105 (.154)	-0.112 (.153)	-0.135 (.157)	-0.118 (.153)
Log Likelihood	-1065.78	-1063.86	-1059.43	-1062.87

model (i) only accounts for opponents' actions;

model (ii) accounts for peers' actions and opponents' responses;

model (iii) is model (ii) with some interaction terms with NX (dummy for sessions with inexperienced subjects);

model (iv) accounts for peers' actions, opponents' responses, and their interaction terms (i.e. outcomes of peers' actions);

^a significantly different from 0 at the 10 percent level

^b significantly different from 0 at the 5 percent level

^c significantly different from 0 at the 1 percent level

standard errors in parentheses

Table 7: SMLEs for the Polya Model (Experiments with Low-Cost Type Entrants)

	model (i)	model (ii)	model (iii)	model (iv)
cumulative entry rate diff. (β)	0.652 (.121) ^c	0.437 (.119) ^c	0.529 (.182) ^b	0.202 (.124)
cumulative entry rate diff. $\times NX$	-	-	-0.155 (.192)	-
weight of entries on peers (ω)	0.055 (.042)	0.014 (.028)	0.013 (.027)	0.010 (.051)
own past choices (λ_1)	0.858 (.150) ^c	0.872 (.161) ^c	1.045 (.178) ^c	0.447 (.183) ^b
own past choices $\times NX$	-	-	-0.253 (.101) ^c	-
own successful limit pricing	-	-	-	0.634 (.122) ^c
depreciation factor (δ_1)	0.634 (.090) ^c	0.598 (.099) ^c	0.590 (.096) ^c	0.587 (.105) ^c
peers' past choices (λ_2)	-	0.964 (.218) ^c	0.617 (.257) ^b	0.861 (.454) ^a
peers' past choices $\times NX$	-	-	0.827 (.432) ^a	-
peers' successful limit pricing	-	-	-	0.235 (.490)
depreciation factor (δ_2)	-	0.784 (.135) ^c	0.822 (.116) ^c	0.793 (.126) ^c
% of M_H choosing 3,4 (cumulative)	0.480 (.209) ^b	0.302 (.207)	-0.275 (.280)	0.313 (.209)
% of M_H choosing 3,4 $\times NX$	-	-	1.248 (.424) ^c	-
experience in a session (γ)	-0.024 (.059)	-0.066 (.059)	-0.101 (.062)	-0.052 (.059)
constant (α)	-1.576 (.113) ^c	-1.613 (.115) ^c	-1.856 (.134) ^c	-1.629 (.118) ^c
random effects (σ)	0.695 (.073) ^c	0.723 (.074) ^c	0.699 (.076) ^c	0.745 (.076) ^c
serial correlation (ρ)	0.251 (.096) ^c	0.223 (.102) ^b	0.233 (.102) ^b	0.208 (.111) ^a
Dummies:				
sessions w/ experienced players	0.735 (.097) ^c	0.574 (.102) ^c	1.139 (.183) ^c	0.605 (.104) ^c
sessions w/ meaningful context	-0.150 (.095)	-0.127 (.096)	-0.117 (.097)	-0.089 (.099)
Log Likelihood	-1647.38	-1635.95	-1624.92	-1621.65

model (i) only accounts for opponents' actions;

model (ii) accounts for peers' actions and opponents' responses;

model (iii) is model (ii) with some interaction terms with NX (dummy for sessions with inexperienced subjects);

model (iv) accounts for peers' actions, opponents' responses, and their interaction terms (i.e. outcomes of peers' actions);

^a significantly different from 0 at the 10 percent level

^b significantly different from 0 at the 5 percent level

^c significantly different from 0 at the 1 percent level

standard errors in parentheses

Table 8: Simulated Percentage of Limit Pricing

	Markov Model			Polya Model	
(a)	(b)	(c)	(b)	(c)	
Experiments with High-Cost Entrants					
<i>Inexperienced subjects</i>					
cycle 1	42.2%	45.4%	44.9%	45.3%	40.3%
cycle 2	58.6%	56.1%	55.2%	57.9%	49.5%
cycle 3	57.6%	60.7%	59.8%	63.3%	53.3%
<i>Experienced subjects</i>					
cycle 1	66.8%	66.4%	65.8%	66.3%	59.5%
cycle 2	63.1%	76.4%	75.4%	78.2%	68.3%
cycle 3	45.8%	79.4%	78.5%	83.2%	73.4%
Experiments with Low-Cost Entrants					
<i>Inexperienced subjects</i>					
cycle 1	14.1%	14.7%	13.6%	15.7%	13.4%
cycle 2	24.0%	20.6%	18.0%	22.9%	17.4%
cycle 3	29.0%	23.6%	20.1%	27.0%	19.0%
<i>Experienced subjects</i>					
cycle 1	41.2%	45.9%	41.1%	45.1%	35.2%
cycle 2	56.3%	67.7%	56.5%	64.9%	46.3%
cycle 3	66.2%	73.4%	59.6%	76.6%	52.3%

(a) % of limit pricing in the raw data;

(b) simulated % of limit pricing with estimated coefficients;

(c) simulated % of limit pricing with estimated coefficients

under the restriction that $\lambda_2 = 0$.

Table 9: Simulated Average Payoff of Monopolists

			Markov Model		Polya Model	
	(a)	(b)	(c)	(d)	(c)	(d)
Experiments with High-Cost Entrants						
<i>Inexperienced subjects</i>						
cycle 1	306.0	308.4	308.3	308.3	308.3	308.5
cycle 2	304.5	299.0	298.7	298.5	298.9	297.8
cycle 3	375.9	326.8	330.3	329.3	333.4	321.8
<i>Experienced subjects</i>						
cycle 1	329.9	314.7	314.6	314.3	314.5	311.4
cycle 2	375.0	322.4	341.3	339.9	343.9	329.8
cycle 3	402.9	327.2	374.1	372.9	379.4	365.7
Experiments with Low-Cost Entrants						
<i>Inexperienced subjects</i>						
cycle 1	498.5	506.7	506.7	506.8	506.6	506.8
cycle 2	516.6	492.4	491.3	490.5	492.0	490.3
cycle 3	513.1	478.6	475.9	474.2	477.6	473.7
<i>Experienced subjects</i>						
cycle 1	555.5	500.6	504.9	500.4	504.1	494.9
cycle 2	564.4	513.4	526.7	513.7	523.5	501.8
cycle 3	564.4	522.2	531.2	514.1	535.3	504.9

(a) expected payoff of limit pricing based on the observed entry rates;

(b) average payoff in the raw data;

(c) simulated average payoff with estimated coefficients;

(d) simulated average payoff with estimated coefficients

under the restriction that $\lambda_2 = 0$.

Figure 1: Experiments with High-Cost Type Entrants under Meaningful Context
 Inexperienced-Subject Sessions

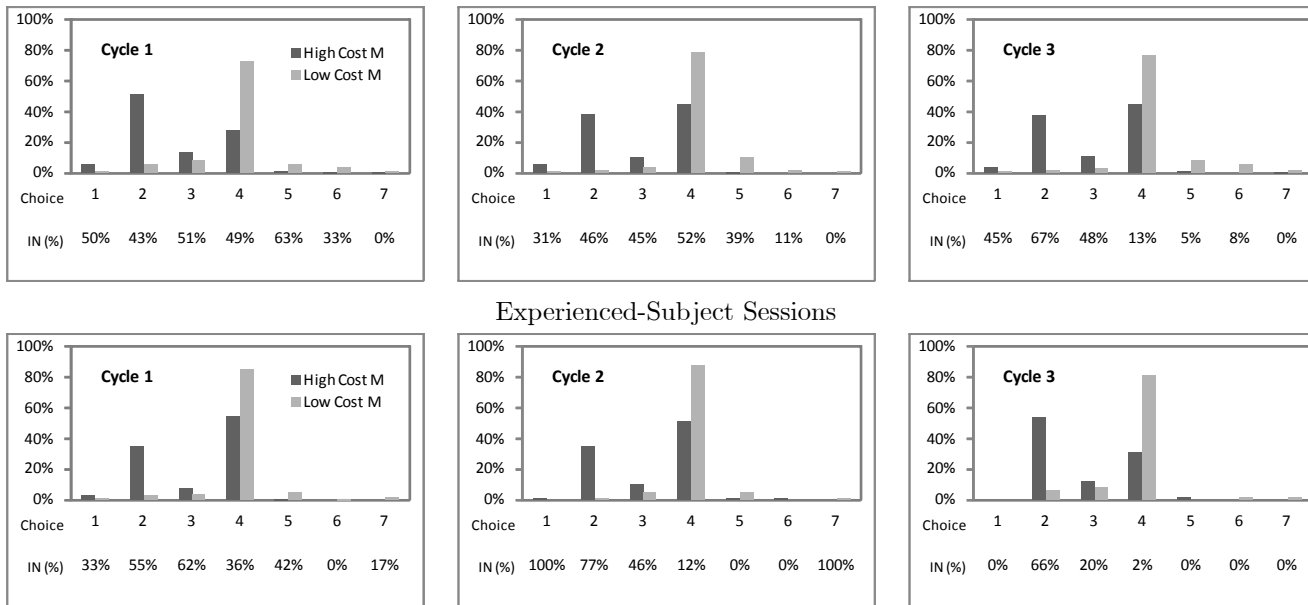



Figure 2: Experiments with Low-Cost Type Entrants under Meaningful Context
 Inexperienced-Subject Sessions



Experienced-Subject Sessions

