When Does Learning in Games Generate Convergence to Nash Equilibria? The Role of Supermodularity in an Experimental Setting

By YAN CHEN AND ROBERT GAZZALE*

This study clarifies the conditions under which learning in games produces convergence to Nash equilibria in practice. We experimentally investigate the role of supermodularity, which is closely related to the more familiar concept of strategic complementarities, in achieving convergence through learning. Using a game from the literature on solutions to externalities, we find that supermodular and "near-supermodular" games converge significantly better than those far below the threshold of supermodularity. From a little below the threshold to the threshold, the improvement is statistically insignificant. Increasing the parameter far beyond the threshold does not significantly improve convergence. (*JEL* C91, D83)

When do players learn to play Nash equilibria? The answer to this important question will help us identify when the outcomes predicted by theory will be realized in competitive environments involving real people. This question has been examined both theoretically (see Drew Fudenberg and David Levine, 1998, for a survey) and experimentally (see Colin Camerer, 2003, for a survey).

According to the theoretical literature, games with strategic complementarities (Paul Milgrom and John Roberts, 1991; Milgrom and Chris Shannon, 1994) have robust dynamic stability properties: under numerous learning dynamics, they converge to the set of Nash equilibria that bound the serially undominated set. The learning dynamics include Bayesian learning, fictitious play, adaptive learning, Cournot best reply, and many others. These games include the supermodular games of Donald Topkis (1979), Xavier Vives (1985, 1990), Russell Cooper and Andrew John (1988), and Milgrom and Roberts (1990). In supermodular games, each player's marginal utility of increasing her strategy rises with increases in her rival's strategies, so that (roughly) the players' strategies are "strategic complements."

Existing literature recognizes that games with strategic complementarities encompass important economic applications of noncooperative game theory, for example, macroeconomics under imperfect competition (Cooper and John, 1988), search (Peter A. Diamond, 1982), bank runs (Douglas W. Diamond and Philip H. Dybvig, 1983; Andrew Postlewaite and Vives, 1987), network and adoption externalities (Dybvig and Chester S. Spatt, 1983), and mechanism design (Chen, 2002).

Past experimental studies of learning and mechanism design suggest that, in addition to equilibrium efficiency, mechanism choice should depend on whether players learn to play the equilibrium and the nature of play on the path to equilibrium. In reviewing the experimental literature on incentive-compatible

* Chen: School of Information, University of Michigan, 1075 Beal Avenue, Ann Arbor, MI 48109 (email: yanchen@umich.edu); Gazzale: Department of Economics, Williams College, Fernald House, Williamstown, MA 01267 (email: rgazzale@williams.edu). We thank Jim Andreoni, Charles Brown, David Cooper, Julie Cullen, John DiNardo, Jacob Goeree, Ernan Haruvy, Daniel Houser, Peter Katuscak, and Jean Krisch for helpful discussions, seminar participants at the University of Michigan, Virginia Tech, and ESA 2001 (Barcelona and Tucson) for their comments, Emre Sucu for programming for the experiment, and Yuri Khoroshilov, Jim Leady, and Belal Sabki for excellent research assistance. We are grateful for the constructive comments of B. Douglas Bernheim, the co-editor, and three anonymous referees, which significantly improved the paper. The research support provided by NSF Grant SES-0079001 to Chen is gratefully acknowledged. Any errors are our own.

mechanisms for pure public goods, Chen (forthcoming) finds that mechanisms with strategic complementarities, such as the Groves-Ledyard mechanism under a high punishment parameter, converge robustly to the efficient equilibrium (Chen and Charles R. Plott, 1996; Chen and Fang-Fang Tang, 1998). Conversely, those far away from the threshold of strategic complementarities do not seem to converge (Vernon L. Smith, 1979; Ronald M. Harstad and Michael Marrese, 1982). In previous experiments, parameters are set either far away from the threshold for strategic complementarities (e.g., Chen and Tang, 1998) or very close to the threshold (e.g., Josef Falkinger et al., 2000).[1] These experiments do not, however, systematically set the parameters below, close to, at, and above the threshold to assess the effect of strategic complementarities on convergence. This is the first such systematic experimental study of games with strategic complementarities.[2] Consequently, this study answers three important questions that the theory on games with strategic complementarities does not address. First, as the parameters approach the threshold of strategic complementarities, will play converge to equilibrium gradually or abruptly? Second, is there a clear performance ranking among games with strategic complementarities? Third, how important is strategic complementarity compared to other factors? The answer to the first question can help us assess, a priori, whether a game "close" to being supermodular, such as the Falkinger mechanism (Falkinger, 1996), might also have good convergence properties. The answer to the second question will help us choose the best parameters within the class of games with strategic complementarities. The answer to the third question will help us assess the importance of strategic complementarities in learning and convergence.

To address these questions, this study adopts an experimental game from the literature on solutions to externalities. Hal R. Varian (1994) proposes a simple class of two-stage mechanisms, the compensation mechanisms, in which subgame-perfect equilibria implement efficient allocations. In a generalized version of Varian's mechanism (John Cheng, 1998), one can vary, without altering the equilibrium, a parameter that determines whether the condition for strategic complementarities is satisfied.

There have been two experimental studies of the compensation mechanisms, neither of which adopts a version with strategic complementarities. James Andreoni and Varian (1999) study the mechanism in the context of the Prisoners' Dilemma. They find that adding a commitment stage to the standard Prisoners' Dilemma game nearly doubles the amount of cooperation to two-thirds. Yasuyo Hamaguchi et al. (2003) investigate a version with a larger strategy space and find 20-percent Nash equilibrium play.

In this paper, we examine the compensation mechanism in an economic environment with a much larger strategy space. Furthermore, we choose various versions to study systematically the effect of strategic complementarities on convergence.

The paper is organized as follows. Section I introduces games with strategic complementarities and presents theoretical properties of the compensation mechanisms. Section II presents the experimental design. Section III introduces the set of hypotheses. Section IV presents experimental results on the level and speed of convergence, as well as efficiency. Section V presents the calibration of three learning models, validation of the models on a hold-out sample, and simulation of performance in the long run using a calibrated learning model. Section VI discusses our findings and Section VII concludes.

I. Strategic Complementarity and the Compensation Mechanisms

In this section, we first introduce games with strategic complementarities and their theoretical properties. We then introduce and analyze a family of mechanisms in this class of games, namely the compensation mechanisms.

Games with strategic complementarities (Milgrom and Shannon, 1994) are those where, given an ordering of strategies, higher actions by one player provide an incentive for the other

[1] Proofs of supermodularity of the Groves-Ledyard and the Falkinger mechanisms are presented in Chen (in press).

[2] By systematically varying the free parameter in the Groves-Ledyard mechanism, Jasmina Arifovic and John O. Ledyard (2001) study learning dynamics and mechanism convergence using genetic algorithms compared with experimental data.

players to take higher actions as well. These games include supermodular games in which the incremental return to any player from increasing her strategy is a nondecreasing function of the strategy choices of other players (*increasing differences*). Furthermore, if a player's strategy space has more than one dimension, components of a player's strategy are complements (*supermodularity*). Membership in this class of games is easy to check. Indeed, for smooth functions in $\mathbb{R}^n$, letting P_i be the strategy space and π_i be the payoff function of player i, the following theorem characterizes increasing differences and supermodularity.

THEOREM 1 (Topkis, 1978): *Let π_i be twice continuously differentiable on P_i. Then π_i has increasing differences in (p_i, p_j) if and only if $\partial^2\pi_i/\partial p_{ih}\partial p_{jl} \geq 0$ for all $i \neq j$ and all $1 \leq h \leq k_i$ and all $1 \leq l \leq k_j$; and π_i is supermodular in p_i if and only if $\partial^2\pi_i/\partial p_{ih}\partial p_{il} \geq 0$ for all i and all $1 \leq h < l \leq k_i$.*

Increasing differences means that an increase in the strategy of player i's rivals raises her marginal utility of playing a high strategy. The supermodularity requirement ensures complementarity among components of a player's strategies and is automatically satisfied in a one-dimensional strategy space. Note that a supermodular game is a game with strategic complementarities, but the converse is not true.

Supermodular games have interesting theoretical properties. In particular, they are robustly stable. Milgrom and Roberts (1990) prove that, in these games, learning algorithms consistent with adaptive learning converge to the set bounded by the largest and the smallest Nash equilibrium strategy profiles. Intuitively, a sequence is consistent with adaptive learning if players "eventually abandon strategies that perform consistently badly in the sense that there exists some other strategy that performs strictly and uniformly better against every combination of what the competitors have played in the not-too-distant past" (Milgrom and Roberts, 1990). This includes numerous learning dynamics, such as Bayesian learning, fictitious play, adaptive learning, and Cournot best reply. While strategic complementarity is sufficient for convergence, it is not a necessary condition. Thus, while games with strategic complementarities ought to converge robustly to the Nash equilibrium, games without strategic complementarities may also converge under specific learning algorithms. Whether these specific learning algorithms are a realistic description of human learning is an empirical question.

While the theory on games with strategic complementarities predicts convergence to equilibrium, it does not address four practical issues. First, as the parameters of a game approach the threshold of strategic complementarities, does play converge gradually or abruptly? Second, is convergence faster further past the threshold? Third, how important is strategic complementarity compared to other features of a game that might also induce convergence to equilibrium? Last, for supermodular games with multiple Nash equilibria, will players learn to coordinate on a particular equilibrium? We choose a game that allows us to answer the first three questions. The fourth question has been addressed by John Van Huyck et al. (1990) and James Cox and Mark Walker (1998).[3]

Specifically, we use the compensation mechanism to study the role of strategic complementarities in learning and convergence to equilibrium play. In the mechanism, each of two players offers to compensate the other for the "costs" of the efficient choice. Assume that when player 1's production equals x, her net profit is $rx - c(x)$, where r is the market price and $c(\cdot)$ is a differentiable, positive, increasing, and convex cost function. Production causes an externality on player 2, whose payoff is $-e(x)$,

[3] Cox and Walker (1998) study whether subjects can learn to play Cournot duopoly strategies in games with two kinds of interior Nash equilibrium. Their type I duopoly has a stable interior Nash equilibrium under Cournot best-reply dynamics and therefore is dominance solvable (Herve Moulin, 1984). Their type II duopoly has an unstable interior Nash equilibrium and two boundary equilibria under Cournot best-reply dynamics, and therefore is not dominance solvable. They found that after a few periods subjects did play stable interior, dominance-solvable equilibria, but they did not play the unstable interior equilibria nor the boundary equilibria. It is interesting to note that these duopoly games are submodular games. Being two-player games, they are also supermodular (Rabah Amir, 1996). Results of Cox and Walker (1998) illustrate the importance of uniqueness together with supermodularity in inducing convergence.

also assumed to be differentiable, positive, increasing, and convex. The mechanism is a two-staged game where the unique subgame-perfect Nash equilibrium induces the Pareto-efficient outcome of x such that $r = e'(x) + c'(x)$. In the first stage (the announcement stage), player 1 announces p_1, a per-unit subsidy to be paid to player 2, while player 2 simultaneously announces p_2, a per-unit tax to be paid by player 1. Announcements are revealed to both players. In the second stage (the production stage), player 1 chooses a production level x. The payoff to player 1 is $\pi_1 = rx - c(x) - p_2x - \alpha(p_1 - p_2)^2$, while the payoff to player 2 is $\pi_2 = p_1x - e(x)$, where $\alpha > 0$ is a free punishment parameter chosen by the designer.

We study a generalized version of the compensation mechanism (Cheng, 1998), which adds a punishment term, $-\beta(p_1 - p_2)^2$, to player 2's payoff function, thus making the payoff functions

$$(1) \quad \pi_1 = rx - c(x) - p_2x - \alpha(p_1 - p_2)^2,$$

$$\pi_2 = p_1x - e(x) - \beta(p_1 - p_2)^2.$$

Using the generalized version, we solve the game by backward induction. In the production stage, player 1 chooses the quantity that solves $\max_x rx - c(x) - p_2x - \alpha(p_1 - p_2)^2$. The first order condition, $r - c'(x) - p_2 = 0$, characterizes the best response in the second stage, $x(p_2)$. In the announcement stage, player 1 solves $\max_{p_1} rx - c(x) - p_2x - \alpha(p_1 - p_2)^2$. The first order condition is

$$(2) \quad \frac{\partial \pi_1}{\partial p_1} = -2\alpha(p_1 - p_2) = 0$$

which yields the best response function for player 1 as $p_1 = p_2$. Player 2 simultaneously solves $\max_{p_2} p_1x(p_2) - e(x(p_2)) - \beta(p_1 - p_2)^2$. The first order condition is

$$(3) \quad \frac{\partial \pi_2}{\partial p_2} = p_1x'(p_2) - e'(x)x'(p_2) + 2\beta(p_1 - p_2) = 0$$

which characterizes player 2's best response function.

The unique subgame-perfect equilibrium has $p_1 = p_2 = p^*$, where p^* is the Pigovian tax which induces the efficient quantity, x^*. As the equilibrium does not depend on the value of β, it holds for the original version where $\beta = 0$. When β is set appropriately, however, the generalized version is a supermodular mechanism. The following proposition characterizes the necessary and sufficient condition for supermodularity.

PROPOSITION 1 (Cheng, 1998): *The generalized version of the compensation mechanism is supermodular if and only if* $\alpha > 0$ *and* $\beta \geq -\frac{1}{2}x'(p_2)$.

The proof is simple. First, as the strategy space is one-dimensional, the supermodularity condition is automatically satisfied. Second, we use Theorem 1 to check for increasing differences. For player 1, from equation (2), we have $\partial^2\pi_1/\partial p_1\partial p_2 = 2\alpha > 0$, while for player 2, from equation (3), we have $\partial^2\pi_2/\partial p_1\partial p_2 = x'(p_2) + 2\beta$. Therefore, $\partial^2\pi_2/\partial p_1\partial p_2 \geq 0$ if and only if $\beta \geq -\frac{1}{2}x'(p_2)$.

To obtain analytical solutions, we use a quadratic cost function $c(x) = cx^2$, where $c > 0$, and a quadratic externality function $e(x) = ex^2$, where $e > 0$. We now summarize the best response functions, equilibrium solutions, and stability analysis in this environment.

PROPOSITION 2: *Quadratic cost and externality functions yield the following characterizations:*

(*i*) *The best response functions for players* 1 *and* 2 *are:*

$$(4) \quad p_1 = p_2;$$

$$(5) \quad p_2 = \frac{\beta - \frac{1}{4c}}{\beta + \frac{e}{4c^2}} p_1 + \frac{\frac{er}{4c^2}}{\beta + \frac{e}{4c^2}}; \textit{ and}$$

$$x = \max\left\{0, \frac{r - p_2}{2c}\right\}.$$

(*ii*) *The subgame-perfect Nash equilibrium is characterized as*

$$(p_1^*, p_2^*, x^*) = \left(\frac{er}{e+c}, \frac{er}{e+c}, \frac{r}{2(e+c)}\right).$$

(*iii*) *If players follow Cournot best-reply dynamics,* (p_1^*, p_2^*) *is a globally asymptotically stable equilibrium of the continuous time dynamical system for any* $\alpha > 0$ *and* $\beta \geq 0$.

(*iv*) *The game is supermodular if and only if* $\alpha > 0$ *and* $\beta \geq 1/4c$.

PROOF:

See Appendix.

The best-response functions presented in part (i) of Proposition 2 reveal interesting incentives. While player 1 has an incentive to always match player 2's price, player 2 has an incentive to match *only* when player 1 plays the equilibrium strategy. Also, at the threshold for strategic complementarity, $\beta = 1/4c$, player 2 has a dominant strategy, $p_2 = er/(e + c) = p_2^*$. Part (iii) of Proposition 2 extends Cheng (1998), who shows that the original version of the compensation mechanism ($\beta = 0$) is globally stable under continuous-time Cournot best-reply dynamics.[4] Cournot best-reply is, however, a relatively poor description of human learning (Richard Boylan and Mahmoud El-Gamal, 1993). Therefore, global stability under Cournot best reply for any $\beta \geq 0$ does not imply equilibrium convergence among human subjects. Part (iv) characterizes the threshold for strategic complementarity in our experiment, a more robust stability criterion than that characterized by part (iii).

An intuition for how strategic complementarities affect the outcome of adaptive learning can be gained from analysis of the best response functions, equations (4) and (5). While player 1's best response function is always upward sloping, player 2's best response function, equation (5), is nondecreasing if and only if $\beta \geq 1/4c$, i.e., when the game is supermodular. Beyond the threshold for strategic complementarity, both best-response functions are upward sloping and they intersect at the equilibrium. It is easy to verify graphically that adaptive learners, e.g., Cournot best reply, will converge to the equilibrium regardless of where they start.

To examine how α, which is unrelated to strategic complementarity, might affect behavior, we observe that when player 1 deviates from the best response by ε, i.e., $p_1 = p_2 + \varepsilon$, his profit loss is $\Delta\pi_1 = -\alpha\varepsilon^2$. This profit loss, which is proportional to the magnitude of α, is the deviation cost for player 1. Based on previous experimental evidence, the incentive to deviate from best response decreases when the deviation cost increases. Chen and Plott (1996) call it the General Incentive Hypothesis, i.e., the error of game theoretic models decreases as the incentive to best-respond increases. We therefore expect that an increase in α improves player 1's convergence to equilibrium. When player 1 plays equilibrium strategy, player 2's best response is to play equilibrium strategy as well. Therefore, we expect that an increase in α might improve player 2's convergence to equilibrium as well. It is not clear, however, whether the α-effects systematically change the effects of the supermodularity parameter β. We rely on experimental data to test the interaction of the α-effects on β-effects.

II. Experimental Design

Our experimental design reflects both theoretical and technical considerations. Specifically, we chose an environment that allows significant latitude in varying the free parameters to better assess the performance of the compensation mechanism around the threshold of strategic complementarity. We describe this environment and the experimental procedures below.

A. *The Economic Environment*

We use the general payoff functions presented in equation (1) with quadratic cost and externality functions to obtain analytical solutions: $c(x) = cx^2$ and $e(x) = ex^2$. We use the following parameters: $c = 1/80$, $e = 1/40$, $r = 24$. From Proposition 2, the subgame-perfect Nash equilibrium is $(p_1^*, p_2^*, x^*) = (16, 16, 320)$ and the threshold for strategic complementarity is $\beta = 1/4c = 20$.

[4] Cheng (1998) also shows that the original mechanism is locally stable under discrete-time Cournot best-reply dynamics.

In the experiment, each player chooses $p_i \in \{0, 1, ..., 40\}$. Without the compensation mechanism, the profit-maximizing production level is $x = r/2c = 960$, three times higher than the efficient level. To reduce the complexity of player 1's problem, we use a grid size of 10 for the quantity and truncate the strategy space, i.e., player 1 chooses $X \in \{0, 1, ..., 50\}$, where $X = x/10$. The payoff functions presented to the subjects are adjusted accordingly. Truncating the strategy space to $X \leq 50$ also reduces the possibility of player 2's bankruptcy. To reduce payoff asymmetry, we give player 1 a lump-sum payment of 250 points each round. Therefore, the equilibrium payoffs under the mechanism for the two players are $\pi_1 = 1530$ and $\pi_2 = 2560$.

The functional forms and specific parameter values are chosen for several reasons. First, equilibrium solutions are integers. Second, equilibrium prices and quantities do not lie in the center of the strategy space, thus avoiding equilibrium convergence as a result of focal points. Third, there is a salient gap between efficiency with and without the mechanism. Without the mechanism, the profit-maximizing production level is $X = 50$, resulting in an efficiency level of 68.4 percent.[5] With the mechanism, the system achieves 100-percent efficiency in equilibrium. Finally, since the threshold for strategic complementarity is $\beta = 20$, there is a large number of integer β values to choose from both below and above the threshold.

To study how strategic complementarity affects equilibrium convergence, we keep $\alpha = 20$, and vary $\beta = 0$, 18, 20, and 40. To study whether α affects convergence, we also keep $\alpha = 10$, and vary $\beta = 0$, 20.

B. *Experimental Procedures*

Our experiment involves 12 players per session—six player 1s (called Red players in the instructions) and six player 2s (Blue players). Each player remains the same type throughout the experiment. At the beginning of each session, subjects randomly draw a PC terminal number. Each then sits in front of the corresponding terminal and is given printed instructions. After the instructions are read aloud, subjects are encouraged to ask questions. The instruction period lasts between 15 and 30 minutes.

Each round a player 1 is randomly matched with a player 2. Subjects are randomly rematched each round to minimize repeated game effects. The random rematching protocol also minimizes the possibility that players collude on a high-subsidy and low-tax outcome.[6] Each session consists of 60 rounds. As we are interested in learning, there are no practice rounds. Each round consists of two stages:

(i) Announcement Stage: Each player simultaneously and independently chooses a price, $p_i \in \{0, 1, ..., 40\}$;

(ii) Production Stage: After (p_1, p_2) are chosen, player 1's computer displays player 2's price and a payoff table showing her payoff for each $X \in \{0, 1, ..., 50\}$. Player 1 then chooses a quantity, X. The server calculates payoffs and sends each player his payoff, the quantity chosen, and the prices submitted by him and his match.

To summarize, each subject knows both payoff functions, the choices made each round by himself and his match, as well as his per-period and cumulative payoffs. At any point, subjects have ready access to all of this information. The mechanism is thus implemented as a game of complete information. We do not know, however, how subjects processed this information, nor do we know their beliefs about the rational-

[5] We compute the efficiency level by using the ratio of total earnings without the mechanism and total earnings with the mechanism. Without the mechanism, at the production level of $X = 50$ (or $x = 500$), total earning is $\pi_1 + \pi_2 = (rx - cx^2) - ex^2 = 2625$. Therefore, the efficiency level is $2625/(1280 + 2560) = 0.684$.

[6] If the players could commit to maximizing joint profits by choosing p_1 and p_2 cooperatively and x noncooperatively, they would choose, by equation (1), $(p_1, p_2, x) = (4(\alpha + \beta)er/[4(\alpha + \beta)(e + c) - 1], r(4e(\alpha + \beta) - 1)/[4(\alpha + \beta)(e + c) - 1], 2r(\alpha + \beta)/[4(\alpha + \beta)(e + c) - 1])$. With our choice of parameters, we get: $p_1 = 24$, $p_2 = 12$ for $\alpha = 20$ and $\beta = 0$; $p_1 = 19.2$, $p_2 = 14.4$ for $\alpha = 20$ and $\beta = 20$; and $p_1 = 18$, $p_2 = 15$ for $\alpha = 20$ and $\beta = 40$; etc. If, in addition, they could choose (p_1, p_2, x) cooperatively, then for our parameter values, we get $(p_1 - p_2, x) = (\min\{480/[3(\alpha + \beta) - 20], 40\}, \min\{960(\alpha + \beta)/[3(\alpha + \beta) - 20], 50\})$.

TABLE 1—FEATURES OF EXPERIMENTAL SESSIONS

Parameters and treatments				Properties	Equilibrium
		α			
		10	20	Supermodular	(p_1^*, p_2^*, X^*)
	0	$\alpha10\beta00$ (4 sessions)	$\alpha20\beta00$ (5 sessions)		
β	18		$\alpha20\beta18$ (4 sessions)	No	(16, 16, 32)
	20	$\alpha10\beta20$ (5 sessions)	$\alpha20\beta20$ (5 sessions)		
	40		$\alpha20\beta40$ (4 sessions)	Yes	(16, 16, 32)

ity of others. Both of these factors introduce uncertainty in the environment.

Table 1 presents features of the experimental sessions, including parameters, number of independent sessions in each treatment, whether the mechanism is supermodular in that treatment, and equilibrium prices and quantities. Overall, 27 independent computerized sessions were conducted in the Research Center for Group Dynamics (RCGD) lab at the University of Michigan from April to July 2001, and in April 2003. We used zTree to program our experiments. Our subjects were students from the University of Michigan. No subject was used in more than one session, yielding 324 subjects. Each session lasted approximately one-and-a-half hours. The exchange rate for all treatments was one dollar for 4,250 points. The average earnings was \$22.82. Data are available from the authors upon request.

III. Hypotheses

Given the design above, we next identify our hypotheses. To do so, we first define and discuss two measures of convergence: the level and speed.[7] In theory, convergence implies that all players play the stage game equilibrium and no deviation is observed. This is not realistic, however, in an experimental setting. Therefore, we define the following measures.

[7] We thank anonymous referees for suggesting this separation and appropriate measures.

DEFINITION 1: *The level of convergence at round t, $L(t)$, is measured by the proportion of Nash equilibrium play in that round. The level of convergence for a block of rounds, $L_b(t_1, t_2)$, measures the average proportion of Nash equilibrium play between rounds t_1 and t_2, i.e., $L_b(t_1, t_2) = \sum_{t=t_1}^{t_2} L(t)/(t_2 - t_1 + 1)$, where $0 \leq t_1 \leq t_2 \leq T$ and T is the total number of rounds.*

We define the level of convergence for both a round and a block of rounds. The block convergence measure smooths out inter-round variation.

Ideally, the speed of convergence should measure how quickly all players converge to equilibrium strategies. In our experimental setting, however, we never observe perfect convergence. We therefore use a more general definition for the speed of convergence.

DEFINITION 2: *For a given level of convergence, $L^* \in (0, 1]$, the speed of convergence is measured by the first round in which the level of convergence reaches L^* and does not subsequently drop below this level, i.e., τ such that $L(t) \geq L^*$ for any $t \geq \tau$.*

We use Definition 2 for analysis of convergence speed in simulations. Due to oscillation in experimental data, Definition 2 is not practical. We use the following observation which enables us to measure convergence speed using estimates for the slope of $L(t)$, $\Delta L(t) \equiv L(t + 1) - L(t)$, obtained in regression analysis. Let $L_y(t)$ be the convergence level of treatment y at time t. We now relate the slope of $L(t)$ and the initial level of convergence $L(1)$ to the speed of convergence.

OBSERVATION 1: *If* $L_1(1) \geq L_2(1)$ *and* $\Delta L_1(t) > \Delta L_2(t)$ *for all* $t \in [1, T - 1]$, *then, given any* $L^* \in (0, 1]$, *the first treatment converges more quickly than the second treatment.*

Based on theories presented in Section I, we now form our hypotheses about the level and speed of convergence. While theories of strategic complementarities do not make any predictions about the speed of convergence, we form our hypotheses based on previous experiments that incidentally address games with strategic complementarities.

HYPOTHESIS 1: *When* $\alpha = 20$, *increasing* β *from* 0 *to* 20 *significantly increases* (*a*) *the level and* (*b*) *the speed of convergence.*

Hypothesis 1 is based on the theoretical prediction that games with strategic complementarities converge to the unique Nash equilibrium, as well as on previous experimental findings that supermodular games perform robustly better than their non-supermodular counterparts.

HYPOTHESIS 2: *When* $\alpha = 20$, *increasing* β *from* 0 *to* 18 *significantly increases* (*a*) *the level and* (*b*) *the speed of convergence.*

Hypothesis 2 is based on the findings of Falkinger et al. (2000) that average play is close to equilibrium when the free parameter is slightly below the supermodular threshold.

HYPOTHESIS 3: *When* $\alpha = 20$, *increasing* β *from* 18 *to* 20 *significantly increases* (*a*) *the level and* (*b*) *the speed of convergence.*

Since we have not found any previous experimental studies that compare the performance of games with strategic complementarities with those near the threshold, Hypothesis 3 is pure speculation.

HYPOTHESIS 4: *When* $\alpha = 20$, *increasing* β *from* 20 *to* 40 *does not significantly increase either* (*a*) *the level or* (*b*) *the speed of convergence.*

Since we have not found any previous experimental studies within the class of games with strategic complementarities, Hypotheses 4 is again our speculation.

HYPOTHESIS 5: *When* $\beta = 0$ *or* 20, *increasing* α *from* 10 *to* 20 *significantly increases* (*a*) *the level and* (*b*) *the speed of convergence.*

HYPOTHESIS 6: *Changing* α *from* 10 *to* 20 *significantly increases the improvement in* (*a*) *the level and* (*b*) *the speed of convergence that results from increasing* β *from* 0 *to* 20.

Hypothesis 5 is based on experimental findings supporting the General Incentive Hypothesis. Hypothesis 6 is our speculation.

Hypotheses 1 through 6 are concerned with only one measure of performance, convergence to equilibrium. Other measures, such as efficiency and budget balance, can be largely derived from convergence patterns. Therefore, although we omit the formal hypotheses, we present results regarding these measures in Sections IV and V.

IV. Experimental Results

In this section, we compare the performance of the mechanism as we vary α and β. At the individual level, we look at both the level and speed of convergence to subgame-perfect equilibrium and near equilibrium in each of the six different treatments. At the aggregate level, we examine the efficiency and budget imbalance generated by each treatment. In the following discussion, we focus on prices rather than on quantity. Recall that player 1's best response in the production stage is uniquely determined by p_2, and that player 1 has all of the information needed to select this best response. In our experiment, deviations in the production stage tend to be small (the average absolute deviation is less than 1 in 23 of 27 sessions) and does not differ significantly among treatments. Therefore, it is not surprising that, in all of our analyses, the results for quantities largely mirror the results for player 2's price.[8]

Recall that subgame perfect Nash equilib-

[8] Tables are available from the authors upon request.

rium for a stage game is $(p_1^*, p_2^*, X^*) = (16, 16, 32)$. Since the strategy space in this experiment is rather large and the payoff function is relatively flat near equilibrium, a small deviation from equilibrium is not very costly. For example, in the $\alpha 20\beta 20$ treatment, a one-unit unilateral deviation from equilibrium prices costs player 1 \$0.005 and player 2 \$0.014. Therefore, we check the ε-equilibrium play by looking at the proportion of price announcements within ± 1 of the equilibrium price, and the quantity announcement within ± 4 of the equilibrium quantity, since a one-unit change in p_2 results in a four-unit best-response quantity change. Therefore, the ε-equilibrium prediction is $(\varepsilon\text{-}p_1^*, \varepsilon\text{-}p_2^*, \varepsilon\text{-}x^*) = (\{15, 16, 17\}, \{15, 16, 17\}, \{28, \ldots, 32, \ldots, 36\})$.

Figures 1 and 2 contain box and whiskers plots for the prices of each treatment for all 60 rounds by players 1 and 2, respectively. The box represents the ranges of the twenty-fifth and seventy-fifth percentiles of prices, while the whiskers extend to the minimum and maximum prices in each round. The horizontal line within each box represents the median price. Compared with the $\beta 00$ treatments, equilibrium price convergence is clearly more pronounced in the supermodular and near-supermodular treatments.

To analyze the performance of the compensation mechanism, we first compare the convergence level achieved in the last 20 rounds of each treatment. Table 2 reports the level of convergence ($L_b(41, 60)$) to subgame-perfect Nash equilibrium (panels A and B) and ε-Nash equilibrium (panels C and D) for each session under each of the six different treatments, as well as the alternative hypotheses and the corresponding p-values of one-tailed permutation tests[9] under the null hypothesis that the convergence levels in the two treatments (column [8]) are the same. While the proportion of ε-equilibrium play is much higher than the proportion of equilibrium play, the results of the permutation tests largely follow similar patterns. We now formally test our hypotheses regarding convergence level. Parts (i) to (iii) of Result 1 present the effects of the degree of strategic complementarity (β-effects). Parts (iv) and (v) present effects due to changes in α (α-effects).

RESULT 1 (Level of Convergence in the Last 20 Rounds):

(*i*) *When* $\alpha = 20$, *increasing* β *from* 0 *to* 18 *and from* 0 *to* 20 *significantly*[10] *increases the level of convergence in* p_1^*, p_2^* *and* $\varepsilon\text{-}p_2^*$;
(*ii*) *When* $\alpha = 20$, *increasing* β *from* 18 *to* 20 *does not change the level of convergence significantly;*
(*iii*) *When* $\alpha = 20$, *increasing* β *from* 20 *to* 40 *does not change the level of convergence significantly;*
(*iv*) *When* $\beta = 0$, *increasing* α *from* 10 *to* 20 *weakly increases the level of convergence in* p_2^*, *and significantly increases the level of convergence in* $\varepsilon\text{-}p_2^*$, *but has no significant effects on* p_1^* *or* $\varepsilon\text{-}p_1^*$;
(*v*) *When* $\beta = 20$, *increasing* α *from* 10 *to* 20 *weakly increases the level of convergence in* p_1^*, *but not in* $\varepsilon\text{-}p_1^*$, p_2^*, *or* $\varepsilon\text{-}p_2^*$.

SUPPORT: The last two columns of Table 2 report the corresponding alternative hypotheses and permutation test results.

By part (i) of Result 1, we accept Hypotheses 1(a) and 2(a). Part (i) also confirms previous experimental findings that supermodular games perform significantly better than those far from the supermodular threshold. Furthermore, near-supermodular games also perform significantly better than those far from the threshold.

By part (ii), however, we reject Hypothesis 3(a). This is the first experimental result that shows that, from a little below the supermodular threshold ($\beta = 18$) to the threshold ($\beta = 20$),

[9] The permutation test, also known as the Fisher randomization test, is a nonparametric version of a difference of two means t-test. (See, e.g., Sidney Siegel and N. John Castellan, 1988.) The idea is simple and intuitive: by pooling all independent observations, the p-value is the exact probability of observing a separation between the two treatments as the one observed when the pooled observations are randomly divided into two equal-sized groups. This test uses all of the information in the sample, and thus has 100-percent power-efficiency. It is among the most powerful of all statistical tests.

[10] When presenting results throughout the paper, we follow the convention that a significance level of 5 percent or less is *significant,* while a significance level between 5 percent and 10 percent is *weakly significant.*

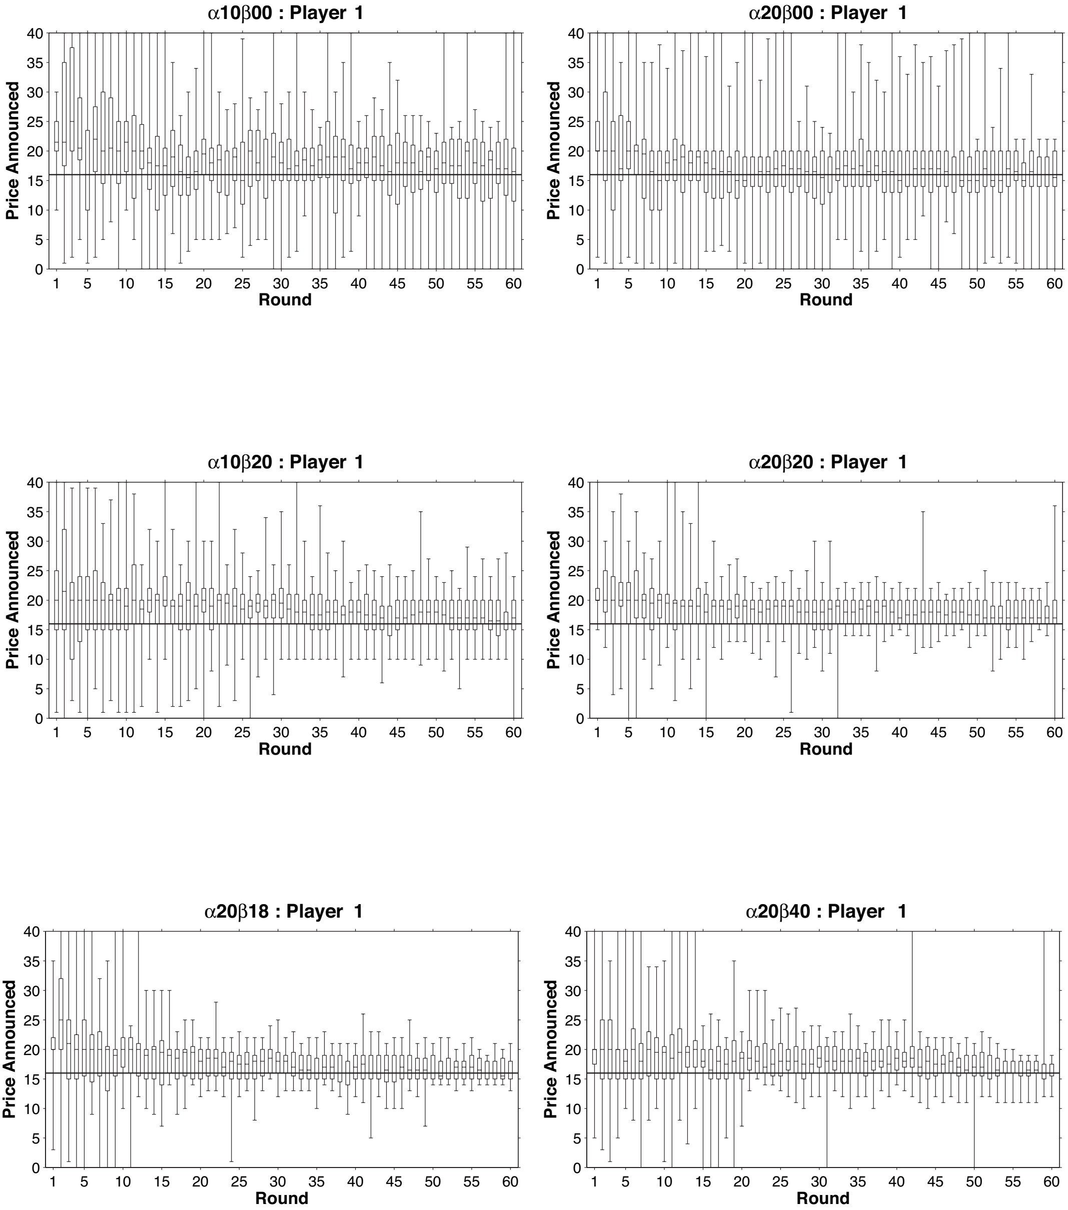

FIGURE 1. DISTRIBUTION OF ANNOUNCED PRICES IN EXPERIMENTAL TREATMENTS: PLAYER 1

improvement in convergence level is statistically insignificant. In other words, we do not see a dramatic improvement at the threshold. This implies that the performance of near-supermodular games, such as the Falkinger mechanism, ought to be comparable to that of supermodular games.

By contrast, we accept Hypothesis 4(a) by part (iii). This is the first experimental result systematically comparing convergence levels of supermodular games, where theory is silent. The convergence level does not significantly improve as β increases from the threshold, 20, to 40. Therefore, the marginal returns for being

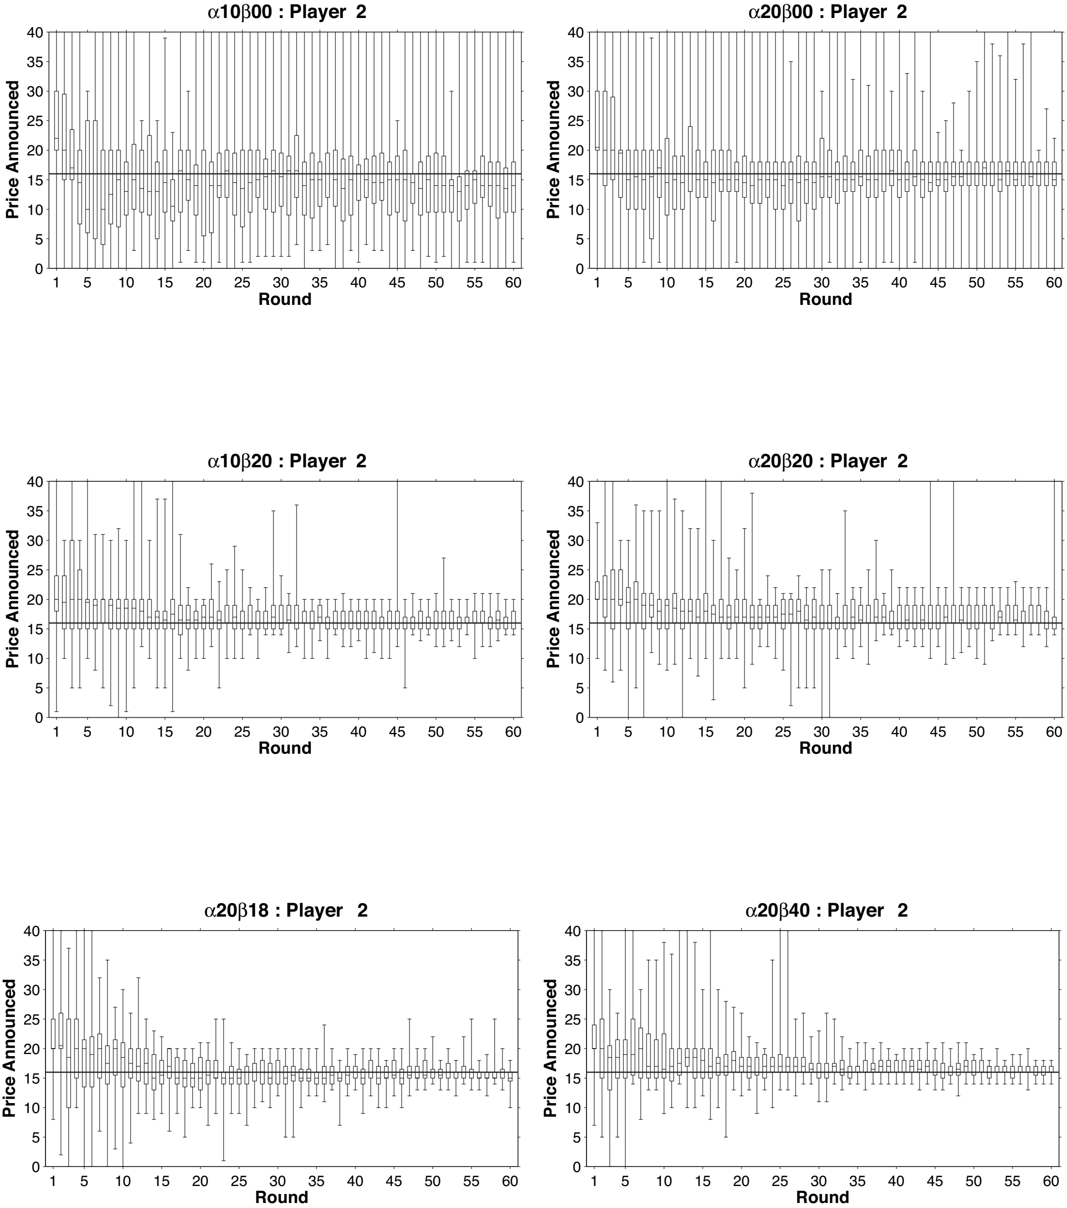

FIGURE 2. DISTRIBUTION OF ANNOUNCED PRICES IN EXPERIMENTAL TREATMENTS: PLAYER 2

"more supermodular" diminish once the payoffs become supermodular.

Proposition 2 predicts convergence under Cournot best reply for any $\beta \geq 0$. However, there is a significant difference in convergence level as β increases from 0 to 18, 20, and beyond. We investigate in Section V whether this difference persists in the long run.

While parts (i) to (iii) present the β-effects, parts (iv) and (v) examine the α-effects and we partially accept Hypothesis 5(a). Recall from equation (5) and subsequent discussions that, at

TABLE 2—LEVEL OF CONVERGENCE IN THE LAST 20 ROUNDS

Panel A:	Proportion of player 1 Nash equilibrium price (p_1^*)						Permutation tests	
Treatment	Session 1	Session 2	Session 3	Session 4	Session 5	Overall	H_1	p-value
$\alpha10\beta00$	0.000	0.008	0.158	0.008		0.044	$\alpha10\beta00 < \alpha20\beta00$	0.397
$\alpha20\beta00$	0.000	0.083	0.083	0.042	0.058	0.053	$\alpha20\beta00 < \alpha20\beta20$	0.008***
$\alpha20\beta18$	0.125	0.117	0.100	0.192		0.133	$\alpha20\beta00 < \alpha20\beta18$	0.008***
$\alpha10\beta20$	0.242	0.067	0.067	0.067	0.208	0.130	$\alpha10\beta20 < \alpha20\beta20$	0.064*
$\alpha20\beta20$	0.325	0.267	0.208	0.058	0.367	0.245	$\alpha20\beta18 < \alpha20\beta20$	0.064*
$\alpha20\beta40$	0.175	0.400	0.158	0.133		0.217	$\alpha20\beta20 < \alpha20\beta40$	0.643

Panel B:	Proportion of player 2 Nash equilibrium price (p_2^*)						Permutation tests	
Treatment	Session 1	Session 2	Session 3	Session 4	Session 5	Overall	H_1	p-value
$\alpha10\beta00$	0.025	0.042	0.025	0.067		0.040	$\alpha10\beta00 < \alpha20\beta00$	0.056*
$\alpha20\beta00$	0.067	0.083	0.033	0.075	0.050	0.062	$\alpha20\beta00 < \alpha20\beta20$	0.032**
$\alpha20\beta18$	0.133	0.292	0.108	0.258		0.198	$\alpha20\beta00 < \alpha20\beta18$	0.008***
$\alpha10\beta20$	0.392	0.108	0.175	0.067	0.667	0.282	$\alpha10\beta20 < \alpha20\beta20$	0.504
$\alpha20\beta20$	0.308	0.117	0.483	0.017	0.450	0.275	$\alpha20\beta18 < \alpha20\beta20$	0.238
$\alpha20\beta40$	0.325	0.492	0.233	0.233		0.321	$\alpha20\beta20 < \alpha20\beta40$	0.365

Panel C:	Proportion of player 1 ε-Nash equilibrium price (ε-p_1^*)						Permutation tests	
Treatment	Session 1	Session 2	Session 3	Session 4	Session 5	Overall	H_1	p-value
$\alpha10\beta00$	0.108	0.200	0.350	0.158		0.204	$\alpha10\beta00 < \alpha20\beta00$	0.183
$\alpha20\beta00$	0.067	0.458	0.358	0.183	0.425	0.298	$\alpha20\beta00 < \alpha20\beta20$	0.087*
$\alpha20\beta18$	0.317	0.625	0.300	0.583		0.456	$\alpha20\beta00 < \alpha20\beta18$	0.103
$\alpha10\beta20$	0.475	0.200	0.300	0.375	0.492	0.368	$\alpha10\beta20 < \alpha20\beta20$	0.194
$\alpha20\beta20$	0.450	0.558	0.475	0.133	0.775	0.478	$\alpha20\beta18 < \alpha20\beta20$	0.444
$\alpha20\beta40$	0.458	0.492	0.375	0.442		0.442	$\alpha20\beta20 < \alpha20\beta40$	0.643

Panel D:	Proportion of player 2 ε-Nash equilibrium price (ε-p_2^*)						Permutation tests	
Treatment	Session 1	Session 2	Session 3	Session 4	Session 5	Overall	H_1	p-value
$\alpha10\beta00$	0.133	0.200	0.242	0.225		0.200	$\alpha10\beta00 < \alpha20\beta00$	0.016**
$\alpha20\beta00$	0.300	0.400	0.200	0.275	0.333	0.302	$\alpha20\beta00 < \alpha20\beta20$	0.016**
$\alpha20\beta18$	0.717	0.667	0.342	0.700		0.606	$\alpha20\beta00 < \alpha20\beta18$	0.016**
$\alpha10\beta20$	0.650	0.517	0.542	0.400	0.892	0.600	$\alpha10\beta20 < \alpha20\beta20$	0.564
$\alpha20\beta20$	0.533	0.592	0.642	0.225	0.900	0.578	$\alpha20\beta18 < \alpha20\beta20$	0.556
$\alpha20\beta40$	0.825	0.792	0.467	0.517		0.650	$\alpha20\beta20 < \alpha20\beta40$	0.294

Note: Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level.

the threshold of strategic complementarity $\beta = 20$, player 2's Nash equilibrium strategy is also a dominant strategy. The finding of no α-effect on player 2's equilibrium play when $\beta = 20$ is consistent with this observation.

To determine the effects of strategic complementarities and other factors on convergence speed, and to investigate further their effects on convergence level, we use probit models with clustering at the individual level. The results from these models are presented in Table 3. The dependent variable is ε-p_1^* in specifications (1) and (3), and ε-p_2^* in specifications (2) and (4).

In specifications (1) and (2), the independent variables are: treatment dummies, D_y, where $y = \alpha10\beta00$, $\alpha20\beta00$, $\beta18$, $\alpha10\beta20$ and $\beta40$; ln(Round); and a constant. We use dummies for different values of α and β rather than direct parameter values to avoid assuming a linear relationship of parameter effects, and omit the dummy for $\alpha20\beta20$. Therefore, in these two specifications, restricting learning speed to be the same across treatments, the estimated coefficient of D_y captures the difference in the convergence level between treatments y and $\alpha20\beta20$. Results from these specifications are

TABLE 3—CONVERGENCE SPEED: PROBIT MODELS WITH CLUSTERING AT INDIVIDUAL LEVEL

	Dependent variable: ε-Nash equilibrium play			
	(1) ε-Nash price 1	(2) ε-Nash price 2	(3) ε-Nash price 1	(4) ε-Nash price 2
$D_{\alpha10\beta00}$	−0.143 (0.048)***	−0.268 (0.050)***		
$D_{\alpha20\beta00}$	−0.090 (0.060)	−0.217 (0.057)***		
$D_{\beta18}$	−0.005 (0.071)	0.004 (0.069)		
$D_{\alpha10\beta20}$	−0.080 (0.060)	0.025 (0.073)		
$D_{\beta40}$	0.000 (0.068)	0.078 (0.078)		
ln(Round)	0.115 (0.014)***	0.167 (0.017)***	0.131 (0.021)***	0.183 (0.022)***
$D_{\alpha10\beta00}$ln(Round)			−0.050 (0.019)***	−0.095 (0.022)***
$D_{\alpha20\beta00}$ln(Round)			−0.031 (0.020)	−0.073 (0.021)***
$D_{\beta18}$ln(Round)			−0.001 (0.020)	0.004 (0.021)
$D_{\alpha10\beta20}$ln(Round)			−0.024 (0.019)	0.008 (0.022)
$D_{\beta40}$ln(Round)			−0.002 (0.019)	0.023 (0.024)
Observations	9720	9720	9720	9720
Number of groups	162	162	162	162
Log pseudo-likelihood	−5562.890	−5824.055	−5557.424	−5793.013
$D_{\alpha10\beta00} = D_{\alpha20\beta00}$	1.50	1.31		
$D_{\alpha10\beta00}$ln(Round) = $D_{\alpha20\beta00}$ln(Round)			1.33	1.23
$D_{\alpha10\beta20} - D_{\alpha10\beta00} = -D_{\alpha20\beta00}$	0.05	1.07		
$D_{\alpha10\beta20}$ln(Round) − $D_{\alpha10\beta00}$ln(Round) = −$D_{\alpha20\beta00}$ln(Round)			0.05	1.03

Notes: Coefficients are probability derivatives. Robust standard errors in parentheses are adjusted for clustering at the individual level. Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level. D_y is the dummy variable for treatment y. Excluded dummy is $D_{\alpha20\beta20}$. The bottom panel presents null hypotheses and Wald $\chi^2(1)$ test statistics.

largely consistent with Result 1 which uses a more conservative test. The coefficients of ln(Round) are both positive and highly significant, indicating that players learn to play equilibrium strategies over time. The concave functional form, ln(Round), which yields a better log-likelihood than either the linear or quadratic functional form, indicates that learning is rapid at the beginning and decreases over time. We will examine learning in more detail in Section V.

In specifications (3) and (4), we use ln(Round), interaction of each of the treatment dummies with ln(Round), and a constant as independent variables. The interaction term allows different slopes for different treatments. Compared with the coefficient of ln(Round), the coefficient for the interaction term, D_y ln(Round), captures the slope differences between treatment y and $\alpha20\beta20$. By Observation 1, as we cannot reject the hypotheses that initial round probability of equilibrium play is the same across all treatments,[11] we can compare the slope of $L(t)$, $\Delta L(t)$, between different treatments. From the probit specification, $L(t) = \Phi(\text{constant} + \mathbf{Db} \ln(t))$, where $\mathbf{D}$ is the 1×5 vector of treatment dummies and $\mathbf{b}$ is the 5×1 vector of estimated coefficients, we can derive

[11] When including treatment dummies in this specification, Wald tests on the hypothesis that the treatment dummy coefficients are all zero yield p-values of 0.2703 for ε-p_1^*, and 0.3383 for ε-p_2^*.

the slope of the probability function for treatment y with respect to t, $\Delta L_y(t) = \Phi'(\text{constant} + \mathbf{Db})b_y/t$. All coefficients reported in Table 3 are increments of probability derivatives, $\Phi'(\text{constant} + \mathbf{Db})b_y$, compared to the baseline case of $\alpha20\beta20$.

RESULT 2 (Speed of Convergence):

(*i*) *When* $\alpha = 20$, *increasing* β *from* 0 *to* 18 *and from* 0 *to* 20 *significantly increases the speed of convergence in* ε-p_1^* *and* ε-p_2^*;
(*ii*) *When* $\alpha = 20$, *increasing* β *from* 18 *to* 20 *has no significant effect on convergence speed;*
(*iii*) *When* $\alpha = 20$, *increasing* β *from* 20 *to* 40 *has no significant effect on convergence speed;*
(*iv*) *When* $\beta = 0$, *increasing* α *from* 10 *to* 20 *has no significant effects on convergence speed;*
(*v*) *When* $\beta = 20$, *increasing* α *from* 10 *to* 20 *has no significant effects on convergence speed.*

SUPPORT: Models (3) and (4) in Table 3 report analyses on convergence speed. For each independent variable, the coefficients, standard errors (in parentheses), and significance levels are reported. The second Wald test looks at whether the coefficient of $D_{\alpha10\beta00}$ln(Round) equals that of $D_{\alpha20\beta00}$ln(Round).

Part (i) of Result 2 supports Hypotheses 1(b) and 2(b), while by part (ii) we reject Hypothesis 3(b). Part (iii) supports Hypothesis 4(b). Parts (iv) and (v) reject Hypothesis 5(b). Result 2 provides the first empirical evidence on the role of strategic complementarity and the speed of convergence.

Although part (ii) indicates that increasing β from 18 to 20 does not significantly change convergence speed, we now investigate whether there are any differences between $\beta18$ and the supermodular treatments. In particular we compare $\alpha20\beta18$ with $\alpha20\beta20$ and $\alpha20\beta40$.[12] In Result 1, we show that these treatments achieve the same convergence level. Also, we cannot reject the hypothesis that round-one prices for each player are drawn for the same distribution for each treatment.[13] As the treatments all start and converge to similar levels of equilibrium play, we use a more flexible functional form to look for differences in speed.

TABLE 4—CONVERGENCE SPEED

Probit models with individual-level clustering comparing $\alpha20\beta18$ with the $\alpha20$ supermodular treatments achieving similar convergence levels

	Dependent variable: ε-Nash equilibrium play	
	(1) ε-Nash price 1	(2) ε-Nash price 2
ln(Round)	0.054	0.216
	(0.040)	(0.043)***
Round	0.004	−0.001
	(0.002)*	(0.002)
$D_{\alpha20\beta18}$ln(Round)	−0.013	−0.066
	(0.032)	(0.035)*
$D_{\alpha20\beta18}$Round	0.001	0.006
	(0.002)	(0.003)**
Observations	4680	4680
Log pseudo-likelihood	−2879.07	−2993.85

Notes: Coefficients reported are probability derivatives. Robust standard errors in parentheses. Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level. D_y is the dummy variable for treatment y. Excluded dummies are $D_{\alpha20\beta20}$ and $D_{\alpha20\beta40}$.

In Table 4, we report the results of probit regressions comparing $\alpha20\beta18$ with the two $\alpha20$ supermodular treatments. We use Round, ln(Round), and their interactions with $\alpha20\beta18$ as independent variables to allow different convergence speeds to the same level of convergence. In model (1), the dependent variable is player 1 ε-equilibrium play. There is no significant difference between $\alpha20\beta18$ and the $\alpha20$ supermodular treatments. In model (2), the dependent variable is player 2 ε-equilibrium play. There are significant differences between the near-supermodular and $\alpha20$ supermodular treatments. In early rounds, ln(Round) is large relative to Round. The negative and weakly significant coefficient on $D_{\alpha20\beta18}$ln(Round) implies a lower probability of early-round equilibrium play in

[12] We omit $\alpha10\beta20$ from this analysis. First, it does not achieve the same ε-p_1^* convergence level as the other supermodular treatments. Second, omitting it avoids the possibility of an α-effect.

[13] Kolmogorov-Smirnov test result tables are available by request.

TABLE 5—EFFICIENCY MEASURE

Efficiency: Last 20 rounds								
Treatment	Session 1	Session 2	Session 3	Session 4	Session 5	Overall	H_1	p-value
$\alpha10\beta00$	0.600	0.671	0.804	0.858		0.733	$\alpha10\beta00 < \alpha20\beta00$	0.087*
$\alpha20\beta00$	0.650	0.870	0.907	0.873	0.825	0.825	$\alpha20\beta00 < \alpha20\beta20$	0.095*
$\alpha20\beta18$	0.963	0.952	0.889	0.959		0.941	$\alpha20\beta00 < \alpha20\beta18$	0.016**
$\alpha10\beta20$	0.958	0.919	0.905	0.881	0.984	0.929	$\alpha10\beta20 < \alpha20\beta20$	0.714
$\alpha20\beta20$	0.888	0.937	0.939	0.780	0.971	0.903	$\alpha20\beta18 < \alpha20\beta20$	0.833
$\alpha20\beta40$	0.973	0.962	0.943	0.949		0.957	$\alpha20\beta20 < \alpha20\beta40$	0.064*

Note: Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level.

$\alpha20\beta18$ relative to the supermodular treatments. The $\beta18$ treatment does catch up to these treatments, which is reflected in the positive and significant coefficient on $D_{\alpha20\beta18}$ interacted with Round. This result suggests a difference between supermodular and near-supermodular treatments: while they achieve similar convergence levels, the supermodular treatments perform better in early rounds and thus might reach certain convergence levels faster than their near-supermodular counterparts.

The previous discussion examines the separate effects of strategic complementarity (β-effects) and α (α-effects) on convergence. Varying α allows us, however, to test the importance of strategic complementarity relative to other features of mechanism design. Having a full factorial design in the parameters $\alpha = 10$, 20 and $\beta = 0$, 20, we can study whether α affects the role of strategic complementarity. In particular, as β increases from 0 to 20, we expect improvement in convergence level and speed. We study whether this improvement changes when α increases from 10 to 20.

The last two Wald tests in the bottom panel of Table 3 separately examine the α-effects on the change in convergence level and speed resulting from increasing β from 0 to 20 (α-effects on β-effects). Changing α from 10 to 20 does not significantly change the improvement in either convergence level or speed. This is the first empirical result examining the effects of other factors on the role of strategic complementarity.

So far, we have discussed the performance of the mechanism relative to equilibrium predictions and individual behavior. We now turn to group-level welfare results. Since the compensation mechanism balances the budget only in equilibrium, total payoffs off the equilibrium path can be only weakly related to efficient payoffs. Therefore, we use two separate measures to capture welfare implications: an efficiency measure and a measure of budget imbalance.

We first define the per-round efficiency measure which includes neither the tax/subsidy nor the penalty terms, i.e.,

$$e(t) = \frac{rx(t) - c(x(t)) - e(x(t))}{rx^* - c(x^*) - e(x^*)}$$

where x^* is the efficient quantity. The efficiency achieved in a block, $e(t_1, t_2)$, where $0 \leq t_1 \leq t_2 \leq 60$, is then defined as

$$e(t_1, t_2) = \sum_{t=t_1}^{t_2} \frac{e(t)}{t_2 - t_1 + 1}.$$

RESULT 3 (Efficiency in the Last 20 Rounds):

(*i*) *When* $\alpha = 20$, *increasing* β *from* 0 *to* 18 *significantly improves efficiency, while increasing* β *from* 0 *to* 20 *weakly improves efficiency;*
(*ii*) *When* $\alpha = 20$, *increasing* β *from* 18 *to* 20 *has no significant effect on efficiency;*
(*iii*) *When* $\alpha = 20$, *increasing* β *from* 20 *to* 40 *weakly improves efficiency;*
(*iv*) *When* $\beta = 0$, *increasing* α *from* 10 *to* 20 *weakly improves efficiency;*
(*v*) *When* $\beta = 20$, *increasing* α *from* 10 *to* 20 *has no significant effect on efficiency.*

SUPPORT: Table 5 presents the efficiency measure for each session in each treatment in the last 20 rounds, the alternative hypotheses, and the results of one-tailed permutation tests.

Result 3 is largely consistent with Result 1, indicating that supermodular and near-supermodular mechanisms induce a significantly higher proportion of equilibrium play than mechanisms far from the supermodular threshold. The new finding is that increasing β from the threshold 20 to 40 weakly improves efficiency.

Second, the budget surplus at round t is the sum of the penalty terms, plus tax, minus subsidy in that round, i.e., $s(t) = (\alpha + \beta)(p_1(t) - p_2(t))^2 + p_2(t)x(t) - p_1(t)x(t)$. Examining the session-level budget surplus across all rounds, we find the following results. First, 22 out of 27 sessions result in a budget surplus. This comes from a combination of our choice of parameters and the dynamics of play. Second, the only significant difference across treatments is that budget surpluses under $\alpha 20\beta 18$ and $\alpha 20\beta 20$ are significantly lower than those under $\alpha 20\beta 40$ (p-values = 0.029 and 0.024 respectively). This finding points to a potential cost of driving up the punishment parameter, β. In other words, before the system equilibrates, a high punishment parameter can worsen the budget imbalance problem inherent in the mechanism.

Overall, Results 1 through 3 suggest the following observations regarding games with strategic complementarities. First, in terms of convergence level and speed, supermodular and near-supermodular games perform significantly better than those far under the threshold. Second, while there is no significant difference between supermodular and near-supermodular games in terms of convergence level, early performance differences may lead to supermodular treatments reaching certain convergence levels more quickly. Third, beyond the threshold, increasing β has no significant effect on either the level or the speed of convergence, but has the disadvantage of risking higher budget imbalance. Last, for a given β, increasing α has partial effects on convergence level, but no significant effect on convergence speed. To check the persistence of these experimental results in the long run, we use simulations in Section V.

V. Simulation Results: Continued Dynamics

Our experiment examines the relationship between strategic complementarity and convergence to equilibrium. Learning theory predicts long-run convergence. Figures 1 and 2 show that convergence continues to improve in later rounds for several treatments, suggesting continued dynamics. Due to time, attention, and resource constraints, it was not feasible to run the experiment for much longer than 60 rounds. Therefore, we rely on simulations to study continued convergence beyond 60 rounds.

To do so, we look for a learning algorithm which, when calibrated, closely approximates the observed dynamic paths over 60 rounds. The large empirical literature on learning in games (see, e.g., Camerer, 2003, for a survey) suggests many models. Our interest here is not to compare the performance of various learning models. We look for learning models that match two criteria. First, we require a model that performs well in a variety of experimental games. Second, given our experiment has complete information about the payoff structure, the model needs to incorporate this information. In the following subsections, we first introduce three learning models which meet our criteria. We then look at how well the learning models predict the experimental data by calibrating each algorithm on a subset of the experimental data and validating the models on a hold-out sample. We then report the forecasting results using one of the calibrated algorithms.

A. *Three Learning Models*

The models we examine are stochastic fictitious play with discounting (hereafter shortened as **sFP**) (Yin-Wong Cheung and Daniel Friedman, 1997; Fudenberg and Levine, 1998), functional EWA (**fEWA**) (Teck-Hua Ho et al., 2001) and the payoff assessment learning model (**PA**) (Rajiv Sarin and Farshid Vahid, 1999). We now give a brief overview of each model. Interested readers are referred to the originals for complete descriptions.

The particular version of sFP that we use is logistic fictitious play (see, e.g., Fudenberg and Levine, 1998). A player predicts the match's price in the next round, $p_{-i}(t + 1)$ according to

$$(6)\quad p_{-i}(t+1) = \frac{p_{-i}(t) + \sum_{\tau=1}^{t-1} r^{\tau} p_{-i}(t-\tau)}{1 + \sum_{\tau=1}^{t-1} r^{\tau}}$$

for some discount factor, $r \in [0, 1]$. Note $r = 0$ corresponds to the Cournot best reply assessment, $p_{-i}(t+1) = p_{-i}(t)$. When $r = 1$, it yields the standard fictitious play assessment. The usual adaptive learning model assumes $0 < r < 1$. All observations influence the expected state but more recent observations have greater weight.

As opposed to standard fictitious play, stochastic fictitious play allows decision randomization and thus better captures the human learning process. Omitting time subscripts, the probability that a player announces price p_i is given by:

$$(7)\quad Pr(p_i|p_{-i}) = \frac{exp(\lambda\pi(p_i, p_{-i}))}{\sum_{j=0}^{40} exp(\lambda\pi(p_j, p_{-j}))}.$$

Given a predicted price, a player is thus more likely to play strategies that yield higher payoffs. How much more likely is determined by λ, the sensitivity parameter. As λ increases, the probability of a best response to p_{-i} increases.

The second model we consider, fEWA, is a one-parameter variant of the experience-weighted attraction (EWA) model (Camerer and Ho, 1999). In this model, strategy probabilities are determined by logit probabilities similar to equation (7) with actual payoffs ($\pi(p_i, p_{-i})$) replaced by strategy *attraction.* In both variants, strategy attraction, $A_i^j(t)$, and an experience weight, $N(t)$, are updated after every period. The experience weight is updated according to $N(t) = \phi(1 - \kappa) \cdot N(t-1) + 1$, where ϕ is the change-detection parameter and κ controls exploration (low κ) versus exploitation. Attractions are updated according to the following rule:

$$(8)\quad A_i^j(t) = \frac{\phi N(t-1)A_i^j(t-1) + [\delta + (1-\delta)I(p_i^j, p_i(t))]\pi_i(p_j, p_{-i}(t))}{N(t)}$$

where the indicator function $I(p_i^j, p_i(t))$ equals one if $p_i^j = p_i(t)$ and zero otherwise, and $\delta \in [0, 1]$ is the imagination weight. In EWA, all parameters are estimated, whereas in fEWA these parameters (except for λ) are endogenously determined by the following functions. The change-detector function, $\phi_i(t)$, is given by

$$(9)\quad \phi_i(t) = 1 - .5\left(\sum_{j=1}^{m_{-i}}\left[\frac{I(p_{-i}^j, p_{-i}(t))}{1} - \frac{\sum_{\tau=1}^{t} I(p_{-i}^j, p_{-i}(t))}{t}\right]^2\right).$$

This function will be close to 1 when recent history resembles previous history. The imagination weight $\delta_i(t)$ equals $\phi_i(t)$, while κ equals the Gini coefficient of previous choice frequencies. EWA models encompass a variety of familiar learning models: cumulative reinforcement learning ($\delta = 0$, $\kappa = 1$, $N(0) = 1$), weighted reinforcement learning ($\delta = 0$, $\kappa = 0$, $N(0) = 1$), weighted fictitious play ($\delta = 1$, $\kappa = 0$), standard fictitious play ($\delta = \phi = 1$, $\kappa = 0$), and Cournot best reply ($\phi = \kappa = 1$, $\delta = 1$).

Finally, we introduce the main components of the PA model. For simplicity, we omit all subscripts that represent player i, and let $\pi_j(t)$ represent the actual payoff of strategy j in round t. Since the game has a large strategy space, we incorporate similarity functions into the model to represent agent use of strategy similarity. As strategies in this game are naturally ordered by their labels, we use the Bartlett similarity function, $f_{jk}(h, t)$, to denote the similarity between the played strategy, k, and an unplayed strategy, j, at period t:

$$f_{jk}(h, t) = \begin{cases} 1 - |j-k|/h & \text{if } |j-k| < h, \\ 0 & \text{otherwise.} \end{cases}$$

In this function, the parameter h determines the $h - 1$ unplayed strategies on either side of the played strategy to be updated. When $h = 1$, $f_{jk}(1, t)$ degenerates into an indicator function equal to one if strategy j is chosen in round t and zero otherwise.

The PA model assumes that a player is a myopic subjective maximizer. That is, he chooses strategies based on assessed payoffs, and does not explicitly take into account the likelihood of alternate states. Let $u_j(t)$ denote the subjective assessment of strategy p_j at time t, and r the discount factor. Payoff assessments are updated through a weighted average of his previous assessments and the payoff he actually obtains at time t. If strategy k is chosen at time t, then:

$$(10)\quad u_j(t+1) = (1 - rf_{jk}(h, t))u_j(t) + rf_{jk}(h, t)\pi_k(t), \qquad \forall j.$$

Each period, the assessed strategy payoffs are subject to zero-mean, symmetrically distributed shocks, $z_j(t)$. The decision maker chooses on the basis of his shock-distorted subjective assessments, $\tilde{u}_j(t) = u_j(t) + z_j(t)$. At time t he chooses strategy p_k if:

$$(11)\quad \tilde{u}_k(t) - \tilde{u}_j(t) > 0, \qquad \forall p_j \neq p_k.$$

Note that mood shocks affect only his choices and not the manner in which assessments are updated.

B. *Calibration*

Literature assessing the performance of learning models contains two approaches to calibration and validation. The first approach calibrates the model on the first t rounds and validates on the remaining rounds. The second approach uses half of the sessions in each treatment to calibrate and the other half to validate. We choose the latter approach for two reasons. First, this approach is feasible as we have multiple independent sessions for each treatment. Second, we need not assume that the parameters of later rounds are the same as those in earlier rounds. We thus calibrate the parameters of each model in blocks of 15 rounds using the experimental data from the first two sessions of each treatment. We then evaluate each model by measuring how well the parameterized model predicts play in the remaining two or three sessions.

For parameter estimation, we conduct Monte Carlo simulations designed to replicate the characteristics of the experimental settings. In all calibrations, we exclude the last two rounds (59 and 60) to avoid any end-of-game effects. We then compare the simulated paths with the experimental data to find those parameters that minimize the mean-squared deviation (MSD) scores. Since the final distributions of our price data are unimodal, the simulated mean is an informative statistic and is well captured by MSD (Ernan Haruvy and Dale Stahl, 2000). In all simulations, we use the k-period-ahead rather than the one-period-ahead approach[14] because we are interested in forecasting the long-run mechanism performance. In doing so, we choose $k = 10$, 15, 20, 30, and 58. We look at blocks of 10, 15, 20, and 30 rounds because as players gain information and experience, information use may change over time. We use $k = 15$ rather than the other values because it best captures the actual dynamics in the experimental data.

Each simulation consists of 1,500 games (18,000 players) and the following steps:

(i) Simulated players are randomly matched into pairs at the beginning of each round;
(ii) Simulated players select price announcements:
  (a) Initial round: Almost half of all players selected a first-round price of 20 (44 percent of player 1s and 43 percent of player 2s). There are also considerable spikes in the frequency of selecting other prices divisible by 5.[15] Based on Kolmogorov-Smirnov tests on the actual round-one price distribution, we reject the null hypotheses of uniform distribution ($d = 0.250$, p-value $= 0.000$) and normal distribution ($d = 0.381$, p-value $= 0.000$). We thus follow the convention (e.g., Ho et al., 2001) and use the actual first-round empirical distribution of choices to generate the first-round choices;
  (b) Subsequent rounds: Simulated player strategies are determined via equation (7) for sFP and fEWA, and equation (11) for PA;

[14] Ido Erev and Haruvy (2000) discuss the tradeoffs of the two approaches.

[15] For example, the seven most frequently selected prices are divisible by 5.

TABLE 6—CALIBRATION OF THREE LEARNING MODELS IN FIFTEEN-ROUND BLOCKS

Model:	sFP				fEWA				PA			
Block:	1	2	3	4	1	2	3	4	1	2	3	4
	Discount rate (r)								**Discount rate** (r)			
$\alpha10\beta00$	1.0	0.7	0.9	1.0					0.1	0.0	0.9	1.0
$\alpha20\beta00$	0.7	0.6	0.1	0.1					0.5	1.0	0.2	0.1
$\alpha20\beta18$	1.0	1.0	0.9	0.9					0.2	1.0	0.3	0.2
$\alpha10\beta20$	1.0	1.0	1.0	0.9					0.1	0.3	0.6	0.3
$\alpha20\beta20$	1.0	1.0	1.0	0.9					0.2	0.1	0.5	0.2
$\alpha20\beta40$	1.0	1.0	1.0	0.7					0.1	1.0	0.4	0.3
	Sensitivity (λ)				**Sensitivity** (λ)				**Shock interval** (a)			
$\alpha10\beta00$	1.5	4.4	4.5	2.8	1.6	4.5	4.4	3.7	500	500	250	50
$\alpha20\beta00$	1.5	4.3	14.0	18.0	1.7	4.0	5.1	7.4	50	100	150	50
$\alpha20\beta18$	1.6	8.0	11.5	18.3	1.5	5.3	6.2	9.5	50	300	500	150
$\alpha10\beta20$	1.9	5.9	8.6	13.8	1.8	4.7	6.2	8.2	100	500	450	300
$\alpha20\beta20$	2.8	5.6	8.1	11.2	2.4	3.8	6.1	6.7	200	300	350	350
$\alpha20\beta40$	3.4	6.5	14.0	20.5	3.2	4.9	6.7	9.9	150	50	150	50
									Similarity window (h)			
$\alpha10\beta00$									1	1	10	9
$\alpha20\beta00$									10	10	9	6
$\alpha20\beta18$									5	4	6	1
$\alpha10\beta20$									1	1	8	3
$\alpha20\beta20$									2	2	7	1
$\alpha20\beta40$									1	3	2	2

(iii) Simulated player 1's quantity choice is based on the following steps:
  (a) Determine each player 1's best response to p_2;
  (b) Determine whether player 1 will deviate from the best response via the actual probability of errors for each block;
  (c) If yes, deviate via the actual mean and standard deviation for the block. Otherwise, play the best response;

(iv) Payoffs are determined by the payoff function of the compensation mechanism, equation (1), for each treatment;

(v) Assessments are updated according to equation (8) for fEWA and equation (10) for PA;

(vi) Proceed to the next round.

The discount factor, $r \in [0, 1]$, is searched at a grid size of 0.1. The parameter λ is searched at a grid size of 0.1 in the interval [1.5, 10.5] for fEWA, and [1.5, 25.0] for sFP. The size of the similarity window, $h \in [1, 10]$, is searched at a grid size of 1. Mood shocks, z, are drawn from a uniform distribution[16] on an interval $[-a, a]$, where a is searched on [0, 500] with a step size of 50. For all parameters, intervals and grid sizes are determined by payoff magnitude.

Table 6 reports the calibrated parameters (discount factor, sensitivity parameter, mood shock interval, and similarity window size) for the first two sessions of each treatment in 15-round blocks.[17] Estimated parameters for the supermodular and near-supermodular treatments are consistent with the increased level of convergence over time, while the $\beta00$ treatments are not. The second column (sFP) reports the best-fit parameters for the stochastic fictitious play model. With the exception of treatment $\alpha20\beta00$, the discount factor is close to 1.0,

[16] Chen and Yuri Khoroshilov (2003) compare three versions of PA models where shocks were drawn from uniform, logistic, and double exponential distributions on two sets of experimental data and find that the performance of the three versions of PA were statistically indistinguishable.

[17] We also calibrate all parameters for the entire 60 rounds, but do not report results due to space limitations.

indicating that players keep track of all past play when forming beliefs about opponents' next move.[18] Given these beliefs, the increasing sensitivity parameter, λ, for all treatments except $\alpha 10\beta 00$ indicates that the likelihood of best response increases over time. The third column reports calibration results for the fEWA model. Again, with the exception of treatment $\alpha 10\beta 00$, the parameter λ increases over time. The final column reports calibration of parameters in the PA model. For each treatment except $\alpha 10\beta 00$, a middle block has the highest discount factor, indicating more weight on new information about a strategy's performance. The second parameter in the PA model represents the upper bound of the interval from which shocks are drawn, a. Experimentation should decrease in the final rounds. Indeed, for all treatments, the estimated mood-shock ranges (weakly) decrease from the third to the last block. Finally, the decreasing similarity windows from the third to final block indicate decreasing strategy spillover. Relatively large discount factors in the third block, combined with relatively large similarity windows, flatten the payoff assessments around the most recently played strategies, consistent with the local experimentation (or relatively stabilized play) observed in the data.

C. *Validation*

Using the parameters calibrated on the first two sessions of each treatment, we next compare the performance of the three learning models in predicting play in the hold-out sample. For comparison, we also present the performance of two static models. The *random choice model* assumes that each player randomly chooses any strategy with equal probability for all rounds. This model incorporates only number of strategies and thus provides a minimum standard for a dynamic learning model. The *equilibrium model* assumes that each player plays the subgame perfect Nash equilibrium every round. Its fit conveys the same information as the proportion of equilibrium play presented in Section IV, but with a different metric (MSD).

Table 7 presents each model's MSD scores for each hold-out session. Recall that two of each treatment's four or five independent sessions are used for calibration, and the rest for validation. The results indicate that all three learning models perform significantly better than the random choice model (p-value < 0.01, one-sided permutation tests) and, by a larger margin, significantly better than the equilibrium model (p-value < 0.01, one-sided permutation tests). While the equilibrium model does a poor job of explaining overall experimental data, its performance improvement over time[19] justifies the use of learning models to explain the dynamics of play. Within each of the top three panels (i.e., learning models), session-level MSD scores are lower for the supermodular and near-supermodular treatments, indicating that each learning model does a better job explaining behavior in those treatments. Indeed, in the empirical learning literature learning models fit better in experiments with better equilibrium convergence (see, e.g., Chen and Khoroshilov, 2003).

RESULT 4 (Comparison of Learning Models): *The performance of PA is weakly better than that of fEWA and strictly better than that of sFP. The performances of fEWA and sFP are not significantly different.*

SUPPORT: Table 7 reports the MSD scores for each independent session in the hold-out sample under each model. The Wilcoxon signed-ranks tests show that the MSD scores under PA are weakly lower than those under fEWA ($z = -0.085$, p-value $= 0.068$), and strictly lower than those under sFP ($z = -0.028$, p-value $=$ 0.023). Using the same test, we cannot reject the hypothesis that fEWA and sFP yield the same MSD scores ($z = 0.662$, p-value $= 0.508$).

Although the performance of PA is only weakly better than that of fEWA, the overall MSD scores for PA are lower than those for

[18] As the discount factor is zero in Cournot best reply, we can reject this model based on our estimation of the discount factors.

[19] We omit the table to show the performance of the equilibrium model in blocks of 15 rounds, as the information is repetitive with Figures 1 and 2.

TABLE 7—VALIDATION ON HOLD-OUT SESSIONS

	Stochastic fictitious play					
Session	$\alpha10\beta00$	$\alpha20\beta00$	$\alpha20\beta18$	$\alpha10\beta20$	$\alpha20\beta20$	$\alpha20\beta40$
1	0.955	0.934	0.940	0.930	0.888	0.940
2	0.960	0.946	0.890	0.917	0.973	0.878
3		0.948		0.883	0.873	
Overall	0.957	0.943	0.915	0.910	0.912	0.909
	fEWA					
1	0.956	0.934	0.942	0.928	0.887	0.948
2	0.960	0.946	0.890	0.922	0.978	0.886
3		0.946		0.879	0.871	
Overall	0.958	0.942	0.916	0.910	0.912	0.917
	Payoff assessment					
1	0.952	0.933	0.914	0.941	0.888	0.916
2	0.957	0.944	0.899	0.922	0.917	0.898
3		0.947		0.894	0.903	
Overall	0.955	0.941	0.907	0.919	0.902	0.907
	Equilibrium play					
1	1.867	1.853	1.833	1.833	1.489	1.792
2	1.889	1.919	1.606	1.839	1.953	1.669
3		1.917		1.447	1.539	
Overall	1.878	1.896	1.719	1.706	1.660	1.731
	Random play					
Overall	0.976	0.976	0.976	0.976	0.976	0.976

fEWA for every treatment. Therefore, we use PA for forecasting beyond 60 rounds.

Figure 3 presents the simulated path of the calibrated PA model and compares it with the actual data in the hold-out sample by superimposing the simulated mean (black line) plus and minus one standard deviation (grey lines) on the actual mean (black box) and standard deviation (error bars). The simulation does a good job of tracking both the mean and the standard deviation of the actual data. In treatments $\alpha10\beta00$ and $\alpha20\beta40$, however, the reduction in variance lags that seen in the middle rounds of the experiment.

D. *Forecasting*

We now report the results from the PA model. As the strategic complementarity predictions concern long-run performance, we use the calibrated parameters for the first 58 rounds to simulate play in later rounds. This exercise allows us to study convergence and efficiency in long but finite horizons.

In all forecasting exercises we base all post-58-round parameters on those from the last block (calibrated from rounds 46 to 58). Since we do not know how these parameters might change beyond 60 rounds, we use three different specifications. First, we retain all final-block parameters. Second, we exponentially decay in subsequent blocks the probability of deviation from the best response in the production stage.[20] Third, we exponentially decay the shock interval in subsequent blocks. As all three specifications yield qualitatively similar results, we report results from only the third specification due to space limitations.[21] In presenting the

[20] In block $k > 4$, we use the probability of deviation, $\text{Pr}_k = \min\{\text{Pr}_4/(k-4)^2, 10^{-8}\}$. We set a lower bound of 10^{-8} to avoid the division-by-zero problem.

[21] Different sequences of random numbers produce slightly different parameters estimates. While the conver-

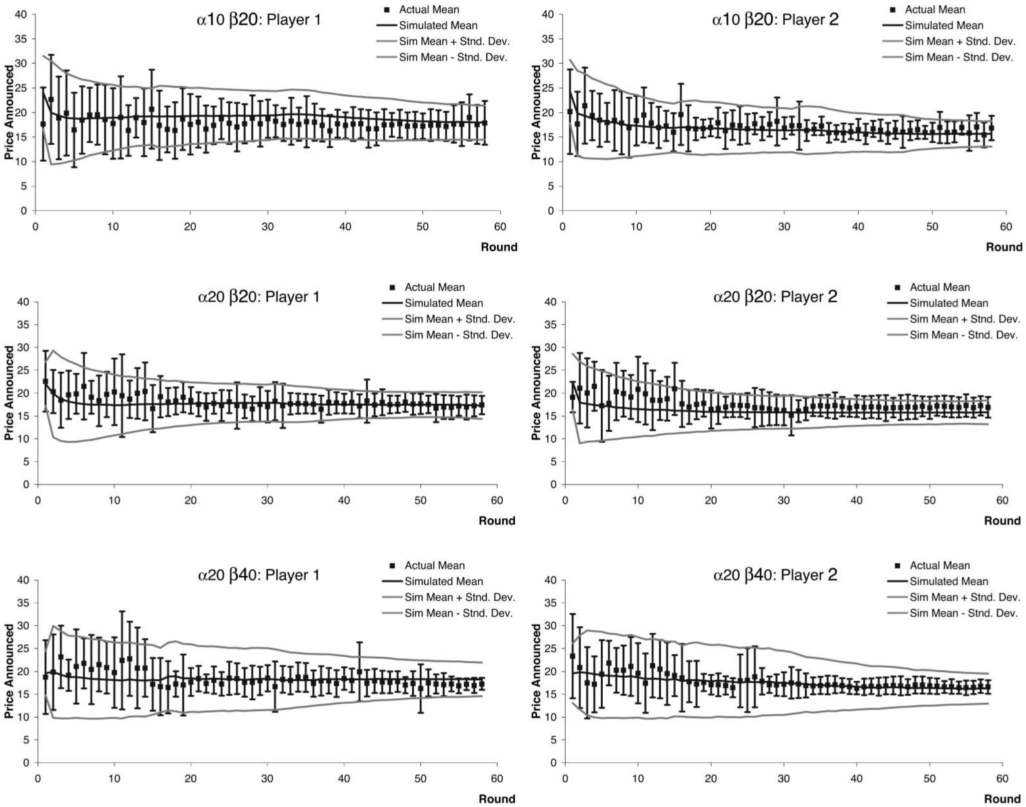

FIGURE 3. SIMULATED DYNAMIC PATH VERSUS ACTUAL DATA

results, we use the shorthand notation $x > y$ to denote a measure under treatment x is significantly higher than that under treatment y at the 5-percent level or less, and $x \sim y$ to denote that the measures are not significantly different at the 5-percent level.

Figure 4 reports the simulated proportion of ε-equilibrium play. We report only the results for the first 500 rounds since the dynamics do not change much thereafter. In our simulation, all treatments reach higher convergence levels than those achieved in the 60 rounds of the experiment. Since the simulated variance reduction lags that of the experiment, round 60 results are not achieved until approximately 90 rounds of simulation. We further note that these improvements slow down after 200 rounds. The proportion of ε-equilibrium play is bounded by 72 percent for player 1 and 93 percent for player 2. We now outline the simulation results.

RESULT 5 (Level of Convergence in Round 500): *In the simulated data for player* 1, *at round* 500, *we have the following level of convergence ranking in:*

(*i*) p_1^*: $\alpha20\beta18 \sim \alpha20\beta40 > \alpha20\beta20 > \alpha10\beta20 > \alpha20\beta00 > \alpha10\beta00$;
(*ii*) ε-p_1^*: $\alpha20\beta18 > \alpha20\beta40 > \alpha20\beta20 > \alpha10\beta20 > \alpha20\beta00 > \alpha10\beta00$;
(*iii*) p_2^*: $\alpha20\beta40 > \alpha20\beta18 \sim \alpha10\beta20 > \alpha20\beta20 > \alpha20\beta00 > \alpha10\beta00$; *and*

gence level of a given treatment for a given set of parameter estimates is not affected by the sequence of random numbers, this convergence level is somewhat sensitive to the exact parameter calibration. We therefore conduct four different calibrations for each treatment and select the parameters that yield the highest level. The relative rankings of treatments are quite robust with respect to ordinal selected.

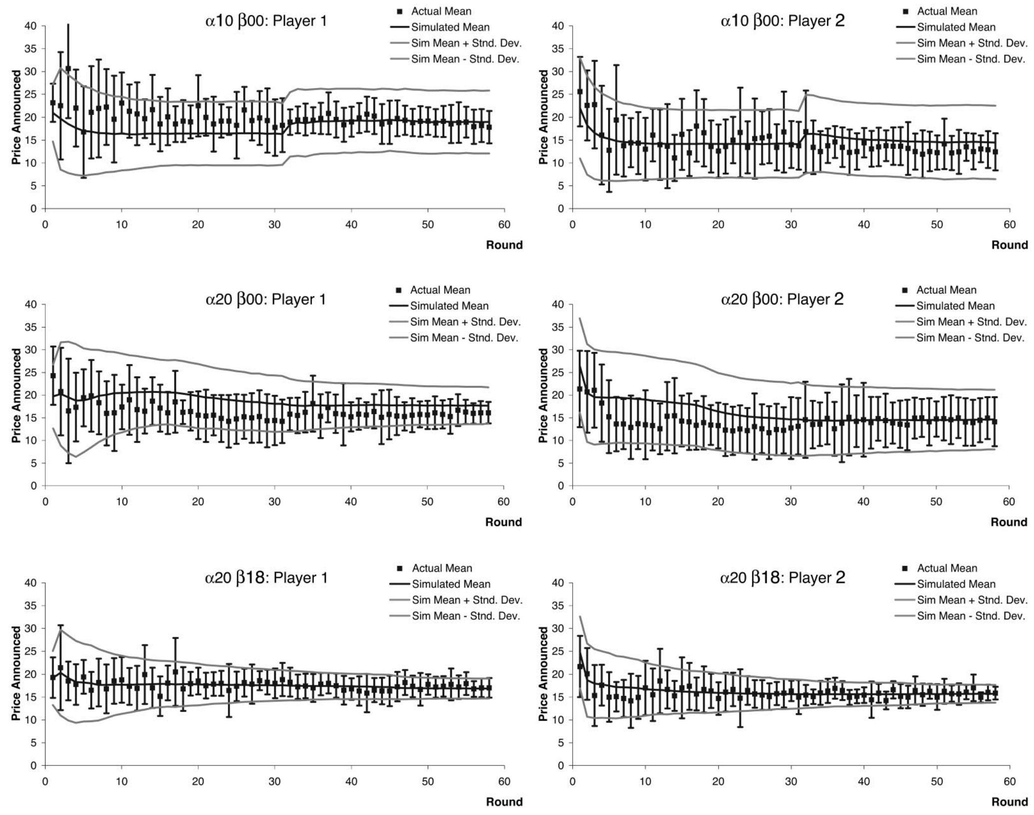

FIGURE 3. CONTINUED

(*iv*) ε-p_2^*: $\alpha20\beta40 > \alpha20\beta18 > \alpha10\beta20 > \alpha20\beta20 > \alpha20\beta00 > \alpha10\beta00$.

SUPPORT: The fourth and seventh columns of Table 8 report p-values for t-tests of the preceding alternative hypotheses for players 1 and 2 respectively.

Comparing Results 1 and 5, we first note that the four supermodular and near-supermodular treatments continue to dominate the $\beta00$ treatments. In fact, the simulations suggest that the gap remains constant compared with $\alpha20\beta00$, and actually increases compared with $\alpha10\beta00$. Unlike Result 1, however, the four dominant treatments differ. In particular, $\alpha20\beta18$ performs better in terms of player 1's convergence, and $\alpha20\beta40$ in terms of player 2's, although the differences within the top four treatments are smaller than the differences between these four and the $\beta00$ treatments. In addition, an increase in α significantly improves convergence by a large margin (40 to 80 percent) for both players when $\beta = 0$.

It is instructive to compare these results with those concerning the speed of convergence in our experimental treatments. In terms of convergence to ε-p_2, our regression results (Table 4) indicate that $\alpha20\beta18$'s performance in later rounds enables it to achieve the same convergence level as the supermodular treatments. This trend continues in our simulations, as $\alpha20\beta18$ performs robustly well compared to the supermodular treatments. Likewise, in our analysis of the experimental results, the convergence speed of $\alpha20\beta40$ was not significantly different from those of the other dominant treatments. In our simulations, however, this treatment dominates all treatments in player 2 equilibrium play.

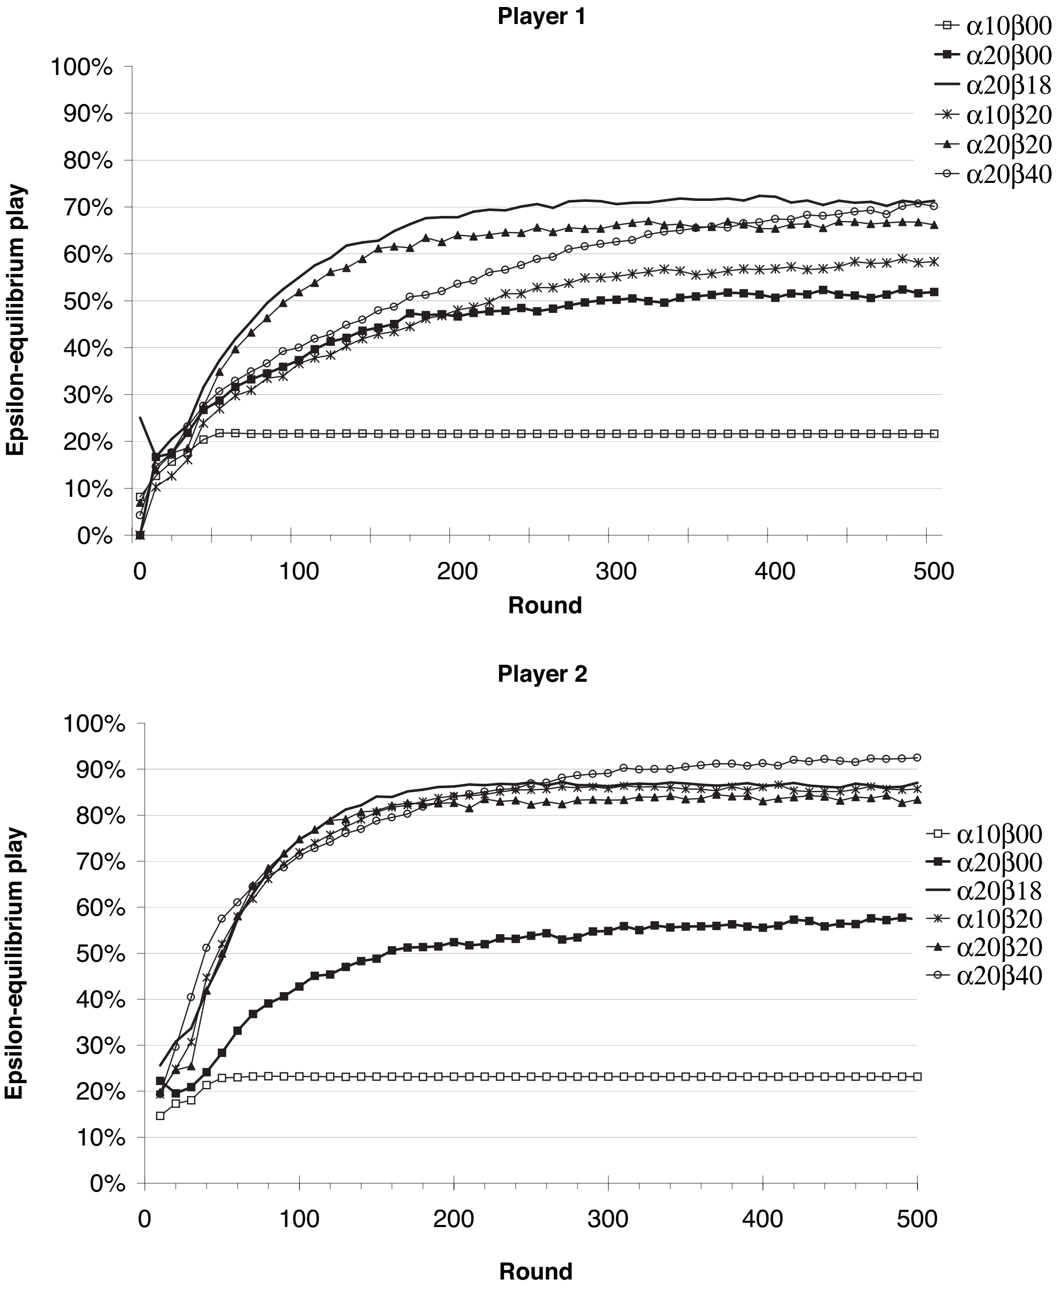

FIGURE 4. SIMULATED ε-EQUILIBRIUM PLAY

In our simulations, the convergence improvement due to increasing β from 0 to 20 does depend on α (α-effect on β-effect). An increase in α significantly decreases the improvement in convergence level (all p-values for one-sided t-tests are less than 0.01). Due to the relatively poor performance of $\alpha 10\beta 00$ in our simulations, increasing α from 10 to 20 when $\beta = 0$

TABLE 8—RESULTS OF t-TESTS COMPARING LEVEL OF CONVERGENCE OF SIMULATIONS OF EXPERIMENTAL TREATMENTS
(*Values are for round 500 and based on 1,500 simulated games.*)

	Player 1 equilibrium price			Player 2 equilibrium price		
Treatment	Probability	H_1	p-value	Probability	H_1	p-value
α10β00	0.073			0.082		
α20β00	0.188	α20β00 > α10β00	0.000***	0.213	α20β00 > α10β00	0.000***
α20β18	0.265	α20β18 > α20β40	0.187	0.381	α20β18 > α10β20	0.264
α10β20	0.211	α10β20 > α20β00	0.000***	0.376	α10β20 > α20β20	0.001***
α20β20	0.243	α20β20 > α10β20	0.000***	0.353	α20β20 > α20β00	0.000***
α20β40	0.259	α20β40 > α20β20	0.007***	0.442	α20β40 > α20β18	0.000***
	Player 1 ε-equilibrium price			**Player 2 ε-equilibrium price**		
Treatment	Probability	H_1	p-value	Probability	H_1	p-value
α10β00	0.216			0.232		
α20β00	0.519	α20β00 > α10β00	0.000***	0.573	α20β00 > α10β00	0.000***
α20β18	0.713	α20β18 > α20β40	0.045**	0.870	α20β18 > α10β20	0.008***
α10β20	0.584	α10β20 > α20β00	0.000***	0.857	α10β20 > α20β20	0.000***
α20β20	0.662	α20β20 > α10β20	0.000***	0.834	α20β20 > α20β00	0.000***
α20β40	0.701	α20β40 > α20β20	0.000***	0.925	α20β40 > α20β18	0.000***

TABLE 9—SPEED OF CONVERGENCE IN SIMULATED DATA
(*Threshold is the percent of players playing epsilon-equilibrium strategy. Entry indicates round in which went over threshold for good, where "—" indicates that treatment never achieved threshold.*)

	Player 1				Player 2					
Threshold	0.4	0.5	0.6	0.7	0.4	0.5	0.6	0.7	0.8	0.9
α20β18	54	83	126	316	38	52	62	87	127	—
α10β20	129	221	—	—	36	48	63	91	148	—
α20β20	61	91	149	—	39	51	61	84	137	—
α20β40	101	166	263	496	30	38	56	94	166	339

improves convergence more dramatically in round 500 than in rounds 40 to 60. In the long run, an increase in α is a partial substitute for an increase in β.

We now use definition 2 to examine the convergence speed in the long run. Table 9 presents results for the first time a treatment reaches the level of convergence, L^*. We omit the two $\beta00$ treatments since they do not converge to the same levels as the other treatments. In terms of player 1, $\alpha20\beta18$ achieves the fastest convergence for all levels L^*. The picture for player 2 convergence, however, is more subtle. First, for the $\beta20$ and $\beta18$ treatments, the speeds of convergence to all L^* are remarkably similar. While $\alpha20\beta40$ lags the other treatments in terms of achieving 80-percent ε-equilibrium play for player 2, it is the only treatment to achieve $L^* = 90$ percent.

Apart from convergence to equilibrium, it is also important to look at long-run welfare properties. We evaluate mechanism performance using the efficiency measure in Section IV.

Figure 5 summarizes the simulated and actual efficiency achieved by each of the treatments. The top panel reports simulated results, while the bottom panel reports efficiency in the experimental data. In the long run, supermodular and near-supermodular treatments continue to differ from the $\beta00$ treatments, with $\alpha20$ superior to $\alpha10$ in the $\beta00$ treatments.

VI. Interpretation and Discussion

The underlying force for convergence in games with strategic complementarities is the combination of the slope of the best-response functions and adaptive learning by players. In

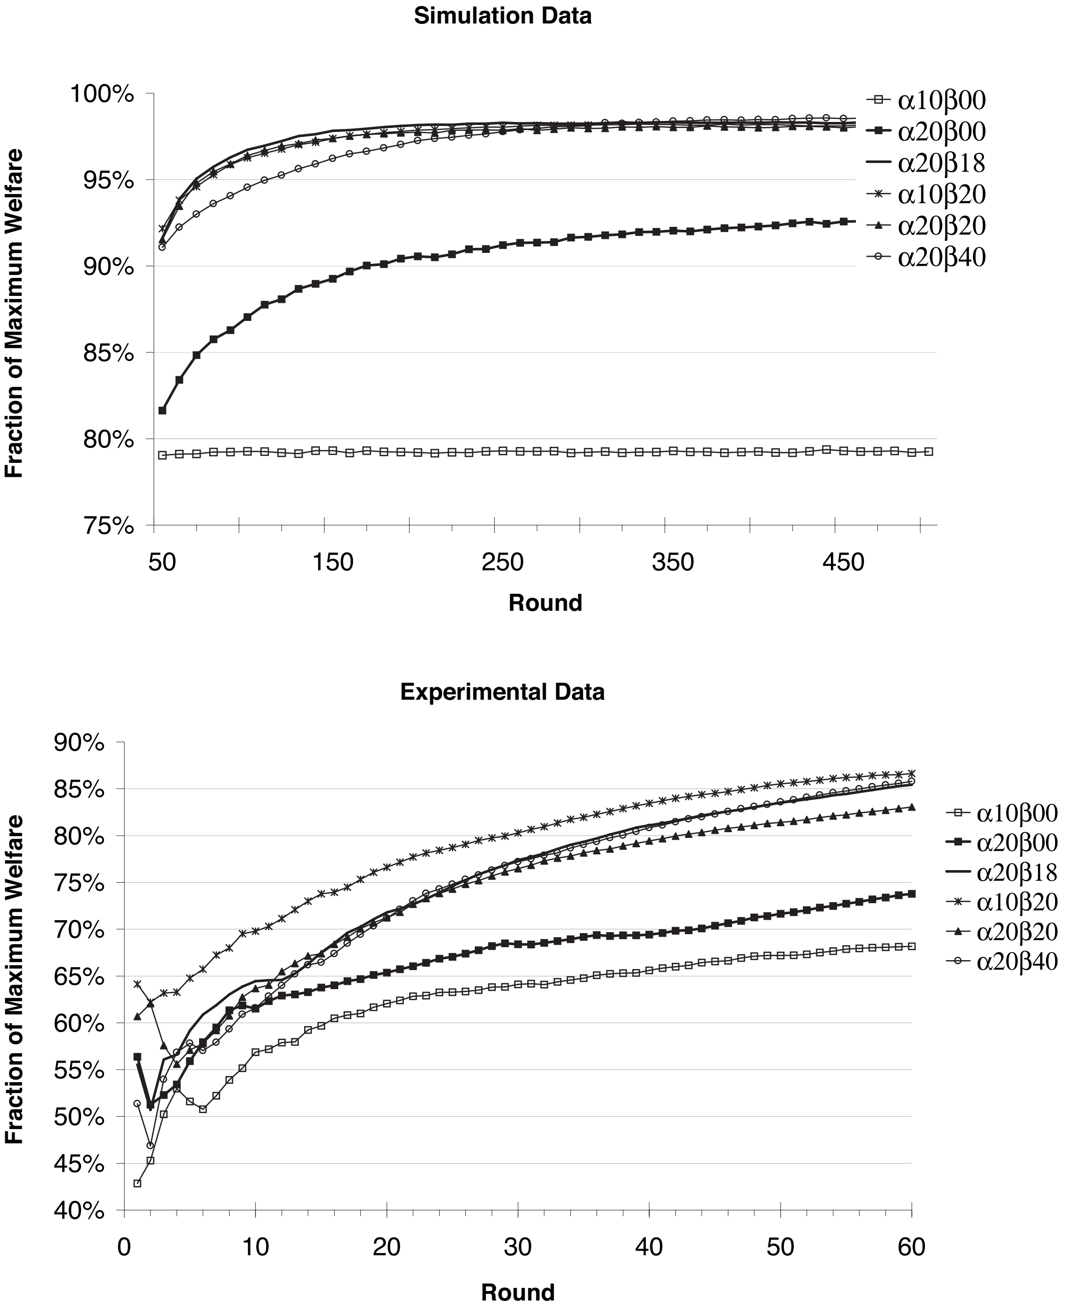

FIGURE 5. EFFICIENCY IN THE SIMULATED AND ACTUAL DATA

the compensation mechanisms, the slope of player 2's best response function is determined by β. As α and β determine the penalty for mismatched prices, we seriously consider the possibility that convergence in supermodular and near-supermodular treatments to an equilibrium where both players announce the same price is not due to strategic complementarities, but rather to increases in α and β making the penalty terms more prominent and matching strategies more focal.[22] We thus have two hypotheses to explain the observed improvement in convergence to equilibrium play in supermodular and near-supermodular treatments: a best-response hypothesis and a matching hypothesis. In this section, we present evidence that the improvement in convergence is due to payoff-relevant changes to best responses, and not to matching becoming more focal.

While the focal point hypothesis suggests that players converge to a match, it does not specify the price at which they match. In the first round of our experiments, almost 50 percent of all prices were 20, and less than 10 percent were in the ε-equilibrium range of [15, 17]. In the final rounds of our experiments, play in the four supermodular or near-supermodular treatments converges strongly to this range, suggesting more than simple matching.

By equation (4), player 1's best response is always to match regardless of p_2. By the General Incentive Hypothesis, we expect more matching behavior by player 1 as α increases. By equation (5), however, matching is a best response for player 2 only in equilibrium. Therefore, data for player 2 can help us separate the best-response and matching hypotheses.

We first investigate which model better explains our experimental data. We operationalize the best response hypothesis by the stochastic fictitious play model of Section V and the matching hypothesis with a stochastic matching model.[23] In both models, the predicted player 1 price is specified by equation (6). In the matching model, player 2 plays this price with probability $1 - \varepsilon$, and one of the other 40 prices with probability $\varepsilon/40$. We estimate ε using the first two sessions of each treatment.[24] We then compare the performance of the matching model with the sFP model using the hold-out sample. Using the Wilcoxon signed-rank tests to compare the MSD scores in the hold-out sample, we conclude that sFP explains player 2's behavior significantly better than the matching model (p-value $= 0.006$). We conclude that best response to a predicted price explains behavior significantly better than the matching model.

We next investigate whether differences in the cost of matching, defined as the difference between matching and best-response profits, can explain all of the differences in matching behavior among treatments, or whether there exists additional matching behavior that cannot be explained by payoff-relevant factors.

We examine the probability that player 2 tries to match the anticipated player 1 price. We do not know what price player 2 anticipates. To estimate how players weigh history, we calibrate a matching model in a manner similar to that outlined in Section V B. We assume a player matches the price he anticipates, given by equation (6), and we find the discount rate for each treatment that minimizes MSD.[25] This gives us the anticipated price that best explains the matching hypothesis. Using this discount rate, we calculate, for each $t > 1$, an anticipated player 1 price, $\bar{p}_1$, for each player 2. We also calculate the opportunity cost of matching $\bar{p}_1$, equal to the difference between best-response and matching profits. This opportunity cost captures the payoff relevant effects of β and also depends on the price to be matched.

We next regress the probability that player 2 matches the anticipated price, $\Pr(p_2 = \bar{p}_1)$, on ln(Round), treatment dummies, and treatment dummies interacted with ln(Round). In one specification, we include the cost of matching as a regressor. If parameter changes make matching more focal, then changes in the cost of matching will not explain all of the changes in probability of matching.

[22] We thank an anonymous referee and the co-editor for pointing this out.

[23] We do not report the results of the deterministic matching model as it performs significantly worse than its stochastic counterpart.

[24] Estimated ε in each block are as follows: $\alpha10\beta00 = \{1.0, .7, .6, .8\}$; $\alpha20\beta00 = \{.7, .6, .5, .4\}$; $\alpha20\beta18 = \{.7, .2, .3, .3\}$; $\alpha10\beta20 = \{.5, .3, .3, .2\}$; $\alpha20\beta20 = \{.1, 0, 0, 0\}$; $\alpha20\beta40 = \{.2, 0, 0, .1\}$.

[25] In all treatments, the calibrated discount rate is $r = 0.1$.

TABLE 10—PLAYER 2 MATCHING HYPOTHESIS
(*Probit models with individual-level clustering comparing models with and without the cost to Player 2 of matching the anticipated price of Player 1.*)

	Dependent variable: Probability player 2 price equals anticipated player 1 price	
Model	(1)	(2)
$D_{\alpha 10}$	0.014	0.017
	(0.043)	(0.041)
$D_{\beta 00}$	−0.077	−0.043
	(0.035)**	(0.036)
$D_{\beta 18}$	−0.046	−0.032
	(0.053)	(0.056)
$D_{\beta 40}$	0.008	−0.041
	(0.061)	(0.042)
ln(Round)	0.041	0.028
	(0.011)***	(0.010)***
Match Cost		−0.001
		(0.000)***
ln(Round)$D_{\alpha 10}$	−0.014	−0.012
	(0.012)	(0.012)
ln(Round)$D_{\beta 00}$	0.003	−0.000
	(0.013)	(0.012)
ln(Round)$D_{\beta 18}$	0.006	0.002
	(0.021)	(0.020)
ln(Round)$D_{\beta 40}$	0.001	0.015
	(0.018)	(0.016)
Observations	9588	9588
Log pseudo-likelihood	−3243.21	−3177.16

Notes: Coefficients reported are probability derivatives. Robust standard errors in parentheses are adjusted for clustering at the individual level. Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level. D_y is the dummy variable for treatment y. Excluded dummy is $D_{\alpha 20\beta 20}$. Match Cost is the difference, in cents, of matching and best-responding to anticipated Player 1 price.

Table 10 presents the results of two probit models with clustering at the individual level. In the first specification, we do not include the cost of matching as a regressor. In this specification, the coefficient for $D_{\beta 00}$ is negative and significant; thus reducing β from 20 to 0 decreases the probability that player 2 will try to match player 1's price. Including the cost of matching in the regression, however, we find that neither the dummies nor their interactions are significant—the coefficient for the previously significant $D_{\beta 00}$ dummy falls almost by a half, while the precision of the estimate stays about the same. The coefficient on the cost of matching, however, is significant and negative. This finding suggests that the rise in matching that occurs when we increase β from 0 to 20 is not because matching is more focal, but rather because an increase of β decreases the cost of matching.[26]

VII. Concluding Remarks

The appeal of games with strategic complementarities is simple: as long as players are adaptive learners, the game will, in the limit, converge to the set bounded by the largest and smallest Nash equilibria. This convergence depends on neither initial conditions nor the assumption of a particular learning model. Unfortunately, while many competitive environments are supermodular, many are not. In this study, we examine whether there are non-supermodular games which have the same convergence properties as supermodular ones. We also study whether there exists a clear convergence ranking among games with strategic complementarities.

Our results confirm the findings of previous experimental studies that supermodular games perform significantly better than games far below the supermodular threshold. From a little below the threshold to the threshold, however, the change in convergence level is statistically insignificant. This results suggests that in the context of mechanism design, the designer need not be overly concerned with setting parameters that are firmly above the supermodular threshold: close is just as good. It also enlarges the set of robustly stable games. For example, the Falkinger mechanism (Falkinger, 1996) is not supermodular, but close. These results suggest that near-supermodular games perform like supermodular ones.

Our next result concerns convergence performance within the class of games with strategic complementarities. Variations in the degree of complementarities have no significant effect on performance within the 60 experimental rounds. Our simulations suggest, however, an increased

[26] We consider other specifications of anticipated price, including the averages of the last one, two, three, and four prices seen. In only one specification was a dummy or its interaction with ln(Round) significant. In that specification, the coefficient for $D_{\beta 18}$ is *positive,* and its interaction with ln(Round) is negative, a finding which is not consistent with a focal point hypothesis.

degree of strategic complementarities leads to improved convergence in the long run.

Finally, the generalized compensation mechanism we use to study convergence has a parameter unrelated to the degree of complementarities. We use this parameter to study the effects factors not related to strategic complementarities on the performance of supermodular games. For a given level of strategic complementarity, the factors have partial effects on convergence level and speed, but the effects of strategic complementarities are largely robust to variations in these other factors. In the long run, these effects persist, and we find evidence that strengthening the strategic complementarities in one player's strategy can partially substitute for the lack of strategic complementarities in the other's.

A word of caution is in order. In a single experimental setting, it is infeasible to study a large number of games in a wide range of complex environments. While this is the first systematic experimental study of the role of strategic complementarities in equilibrium convergence, the applicability of our results to other games needs to be verified in future experiments. In the only other study of this kind, Arifovic and Ledyard (2001) examine similar questions using the Groves-Ledyard mechanism. Their results are encouraging, as their results are consistent with ours.

Appendix: Proof of Proposition 2

We omit the proofs of parts (i), (ii), and (iv), as they follow directly from the previous analysis and Proposition 1. We now present the proof for part (iii). If players follow Cournot best-reply dynamics, we can rewrite equations (4) and (5) as

$$p_1(t+1) = p_2(t);$$

$$p_2(t+1) = mp_1(t) + n$$

where $m = [\beta - (1/4c)]/[\beta + (e/4c^2)]$ and $n = (er/4c^2)/[\beta + (e/4c^2)]$. The analogous differential equation system is

$$\dot{p}_1 = p_2 - p_1; \qquad (12)$$

$$\dot{p}_2 = mp_1 - p_2 + n.$$

We now use the Lyapunov second method to show that system (12) is globally asymptotically stable. Define

$$V(p_1, p_2) \equiv \frac{(p_1 - p_2)^2}{2} + \int_{p_0}^{p_1} (p - mp - n)\, dp$$

where $p_0 \geq 0$ is a constant. We now show that $V(p_1, p_2)$ is the Lyapunov function of system (12).

Define $G(p_1) \equiv \int_{p_0}^{p_1} (p - mp - n)\, dp$. We get

$$G(p_1) = \tfrac{1}{2}(1 - m)p_1^2 - np_1 - \tfrac{1}{2}(1 - m)p_0^2 + np_0.$$

As $c > 0$ and $e > 0$, for any $\beta \geq 0$, we always have $m < 1$. Therefore, the function $G(p_1)$ has a global minimum, which is characterized by the following first-order condition:

$$\frac{dG(p_1)}{dp_1} = (1 - m)p_1 - n = 0.$$

Therefore, $p_1 = n/(1 - m) = er/(e + c) \equiv p^*$ is the global minimum. When $p_1 = p_2 = p^*$, $(p_1 - p_2)^2/2$ also reaches its global minimum. Therefore, $V(p_1, p_2)$ is at its global minimum when $p_1 = p_2 = p^*$.

Next we show that $\dot{V} \leq 0$.

$$\begin{aligned}\dot{V} &= \frac{\partial V}{\partial p_1}\dot{p}_1 + \frac{\partial V}{\partial p_2}\dot{p}_2 \\ &= (p_1 - p_2 + (1 - m)p_1 - n)(p_2 - p_1) \\ &\quad + (p_2 - p_1)(mp_1 - p_2 + n) \\ &= -2(p_1 - p_2)^2 \leq 0.\end{aligned}$$

Let B be an open ball around (p^*, p^*) in the plane. For all $(p_1, p_2) \neq (p^*, p^*)$, $\dot{V} < 0$.

Therefore, (p^*, p^*) is a globally asymptotically stable equilibrium of (12).

REFERENCES

Amir, Rabah. "Cournot Oligopoly and the Theory of Supermodular Games." *Games and Economic Behavior*, 1996, *15*(2), pp. 132–48.

Andreoni, James and Varian, Hal R. "Pre-Play Contracting in the Prisoners' Dilemma." *Proceedings of the National Academy of Science*, 1999, *96*(19), pp. 10933–38.

Arifovic, Jasmina and Ledyard, John O. "Computer Testbeds and Mechanism Design: Application to the Class of Groves-Ledyard Mechanisms for Provision of Public Goods." Unpublished Paper, 2001.

Boylan, Richard T. and El-Gamal, Mahmoud A. "Fictitious Play: A Statistical Study of Multiple Economic Experiments." *Games and Economic Behavior*, 1993, *5*(2), pp. 205–22.

Camerer, Colin F. *Behavioral game theory: Experiments in strategic interaction* (Roundtable Series in Behavioral Economics). Princeton: Princeton University Press, 2003.

Camerer, Colin F. and Ho, Teck-Hua. "Experience-Weighted Attraction Learning in Normal Form Games." *Econometrica*, 1999, *67*(4), pp. 827–74.

Chen, Yan. "A Family of Supermodular Nash Mechanisms Implementing Lindahl Allocations." *Economic Theory*, 2002, *19*(4), pp. 773–90.

Chen, Yan. "Incentive-Compatible Mechanisms for Pure Public Goods: A Survey of Experimental Research," in C. R. Plott and V. L. Smith, eds., *The handbook of experimental economics results*. Amsterdam: Elsevier, forthcoming.

Chen, Yan. "Dynamic Stability of Nash-Efficient Public Goods Mechanisms: Reconciling Theory and Experiments," in Rami Zwick and Amnon Rapoport, eds., *Experimental business research, Vol. 2*. Norwell and Dordrecht: Kluwer Academic Publishers, in press.

Chen, Yan and Plott, Charles R. "The Groves-Ledyard Mechanism: An Experimental Study of Institutional Design." *Journal of Public Economics*, 1996, *59*(3), pp. 335–64.

Chen, Yan and Tang, Fang-Fang. "Learning and Incentive-Compatible Mechanisms for Public Goods Provision: An Experimental Study." *Journal of Political Economy*, 1998, *106*(3), pp. 633–62.

Chen, Yan and Khoroshilov, Yuri. "Learning under Limited Information." *Games and Economic Behavior*, 2003, *44*(1), pp. 1–25.

Cheng, Qiang John. "Essays on Designing Economic Mechanisms." Unpublished Paper, 1998.

Cheung, Yin-Wong and Friedman, Daniel. "Individual Learning in Normal Form Games: Some Laboratory Results." *Games and Economic Behavior*, 1997, *19*(1), pp. 46–76.

Cooper, Russell and John, Andrew. "Coordinating Coordination Failures in Keynesian Models." *Quarterly Journal of Economics*, 1988, *103*(3), pp. 441–63.

Cox, James C. and Walker, Mark. "Learning to Play Cournot Duopoly Strategies." *Journal of Economic Behavior and Organization*, 1998, *36*(2), pp. 141–61.

Diamond, Douglas W. and Dybvig, Philip H. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy*, 1983, *91*(3), pp. 401–19.

Diamond, Peter A. "Aggregate Demand Management in Search Equilibrium." *The Journal of Political Economy*, 1982, *90*(5), pp. 881–94.

Dybvig, Philip H. and Spatt, Chester S. "Adoption Externalities as Public Goods." *Journal of Public Economics*, 1983, *20*(2), pp. 231–47.

Erev, Ido and Haruvy, Ernan. "On the Potential Value and Current Limitations of Data-Driven Learning Models." Unpublished Paper, 2000.

Falkinger, Josef. "Efficient Private Provision of Public Goods by Rewarding Deviations from Average." *Journal of Public Economics*, 1996, *62*(3), pp. 413–22.

Falkinger, Josef; Fehr, Ernst; Gächter, Simon and Winter-Ebmer, Rudlof. "A Simple Mechanism for the Efficient Provision of Public Goods: Experimental Evidence." *American Economic Review*, 2000, *90*(1), pp. 247–64.

Fudenberg, Drew and Levine, David K. "The Theory of Learning in Games," in *MIT Press series on economic learning and social evo-*

lution, Vol. 2. Cambridge, MA: MIT Press, 1998.

Hamaguchi, Yasuyo; Maitani, Satoshi and Saijo, Tatsuyoshi. "Does the Varian Mechanism Work? Emissions Trading as an Example." *International Journal of Business and Economics*, 2003, *2*(2), pp. 85–96.

Harstad, Ronald M. and Marrese, Michael. "Behavioral Explanations of Efficient Public Good Allocations." *Journal of Public Economics*, 1982, *19*(3), pp. 367–83.

Haruvy, Ernan and Stahl, Dale. "Initial Conditions, Adaptive Dynamics, and Final Outcomes: An Approach to Equilibrium Selection." Unpublished Paper, 2000.

Ho, Teck-Hua; Camerer, Colin F. and Chong, Juin-Kuan. *Economic value of EWA Lite: A functional theory of learning in games*. California Institute of Technology, Division of the Humanities and Social Sciences, Working Paper No. 1122-2001.

Milgrom, Paul R. and Shannon, Chris. "Monotone Comparative Statics." *Econometrica*, 1994, *62*(1), pp. 157–80.

Milgrom, Paul R. and Roberts, John. "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities." *Econometrica*, 1990, *58*(6), pp. 1255–77.

Milgrom, Paul R. and Roberts, John. "Adaptive and Sophisticated Learning in Normal Form Games." *Games and Economic Behavior*, 1991, *3*(1), pp. 82–100.

Moulin, Herve. "Dominance Solvability and Cournot Stability." *Mathematical Social Sciences*, 1984, *7*(1), pp. 83–102.

Postlewaite, Andrew and Vives, Xavier. "Bank Runs as an Equilibrium Phenomenon." *Journal of Political Economy*, 1987, *95*(3), pp. 485–91.

Sarin, Rajiv and Vahid, Farshid. "Payoff Assessments without Probabilities: A Simple Dynamic Model of Choice." *Games and Economic Behavior*, 1999, *28*(2), pp. 294–309.

Siegel, Disney and Castellan, N. John, Jr. *Nonparametric statistics for the behavioral sciences, 2nd ed*. New York: McGraw-Hill, 1988.

Smith, Vernon L. "An Experimental Comparison of Three Public Good Decision Mechanisms." *Scandinavian Journal of Economics*, 1979, *81*(2), pp. 198–215.

Topkis, Donald. "Minimizing a Submodular Function on a Lattice." *Operations Research*, 1978, *26*(2), pp. 305–21.

Topkis, Donald. "Equilibrium Points in Nonzero-Sum N-Person Submodular Games." *SIAM Journal on Control and Optimization*, 1979, *17*(6), pp. 773–87.

Van Huyck, John B.; Battalio, Raymond C. and Beil, Richard O. "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review*, 1990, *80*(1), pp. 234–48.

Varian, Hal R. "A Solution to the Problem of Externalities When Agents Are Well Informed." *American Economic Review*, 1994, *84*(5), pp. 1278–93.

Vives, Xavier. "Nash Equilibrium in Oligopoly Games with Monotone Best Response." University of Pennsylvania, CARESS Working Paper 85-10, 1985.

Vives, Xavier. "Nash Equilibrium with Strategic Complementarities." *Journal of Mathematical Economics*, 1990, *19*(3), pp. 305–21.