What Price Coordination?
The Efficiency-Enhancing Effect of Auctioning the Right to Play

By VINCENT CRAWFORD AND BRUNO BROSETA*

A model is proposed to explain the results of recent experiments in which subjects repeatedly played a coordination game, with the right to play auctioned each period in a larger group. Subjects bid the market-clearing price to a level recoverable only in the efficient equilibrium and then converged to that equilibrium, although subjects playing the game without auctions converged to inefficient equilibria. The efficiency-enhancing effect of auctions is reminiscent of forward induction, but is not explained by equilibrium refinements. The model explains it by showing how strategic uncertainty interacts with history-dependent learning dynamics to determine equilibrium selection. (*JEL* C73, C92, C51)

Coordination is central to many questions in economics, from the determination of bargaining outcomes to the design of incentive schemes, the efficacy of implicit contracts, and the influence of expectations in macroeconomics. Such questions are usually modeled as noncooperative games with multiple Nash equilibria, and analyzed under the assumption that players can realize any desired equilibrium. Yet this begs the questions of whether, and how, coordination comes about and how the environment influences equilibrium selection—questions that lie at the heart of most applications.

Although recent advances in game theory have added much to our understanding of coordination, there remains a large gap between theory and experience that is unlikely to be closed by theory alone. Further progress seems likely to depend on combining theory with empirical evidence on strategic behavior. Experiments are a particularly useful source of such evidence, in part because they make it possible to observe the entire coordination process. Crawford (1997) surveys a number of recent studies in which subjects repeatedly played coordination games, uncertain only about each other's strategy choices. The typical result was convergence to an equilibrium, often with a systematic pattern of equilibrium selection in the limit. Explaining such patterns promises to shed considerable light on coordination, in the field as well as the laboratory.

In this paper we seek to explain some of the most intriguing evidence on coordination we have seen, from the experiment of John Van Huyck et al. (henceforth "VHBB") (1993). Their subjects repeatedly played a nine-person coordination game with seven symmetric Pareto-ranked equilibria from the experiments of VHBB (1991), with subjects' payoffs and best responses determined by their own actions, called *efforts,* and an order statistic of all nine subjects' efforts, in this case the median.[1]

* Crawford: Department of Economics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093; Broseta: Department of Economics, University of Arizona, McClelland Hall, Room 401, Tucson, AZ 85721. We thank Raymond Battalio, Andreas Blume, Miguel Costa Gomes, Yong-Gwan Kim, Mark Machina, Daniel McFadden, Ronald Oaxaca, Alvin Roth, Jason Shachat, Joel Sobel, Glenn Sueyoshi, John Van Huyck, and two referees for helpful comments; Miguel Costa Gomes and Mahesh Johari for research assistance; Raymond Battalio, Richard Beil, Gerard Cachon, Colin Camerer, and John Van Huyck for sharing their experimental data; Raymond Battalio, John Van Huyck, and Roberto Weber for conducting additional experiments for us; and the National Science Foundation, the Centro de Formación del Banco de España, and the College of Business and Public Administration, University of Arizona, for research support.

[1] See also VHBB's (1990) experiments, where the order statistic was the minimum.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

However, instead of endowing nine subjects with the right to play the game as in their 1991 experiments, VHBB auctioned its nine positions each period in a group of 18. The structure was publicly announced at the start; the market-clearing price and median effort were publicly announced each period; and explicit communication was prohibited throughout.

VHBB's (1993) design is of unusual economic interest. Its 1991 precursor, in which nine subjects played the same median coordination game without auctions, is among the simplest models of the emergence of conventions to solve coordination problems, with a range of possible outcomes and a natural measure of their efficiency. The median game is similar in structure to several important models, from the Stag Hunt example Rousseau used to motivate his analysis of the social contract to Keynes's beauty contest analogy and the more prosaic macroeconomic coordination models surveyed in Russell Cooper and Andrew John (1988). Its efficient equilibrium is plainly the "correct" coordinating principle, but it is best for a subject to play his part of that equilibrium only if he thinks it likely that enough other subjects will do so. As in the other coordination games VHBB (1990, 1991) studied, there is a tension between the higher payoff of the efficient equilibrium and the greater robustness of other equilibria to subjects' uncertainty about each other's responses, which we call *strategic uncertainty*. Finally, auctioning the right to play is an interesting form of preplay communication, in which subjects' willingness to pay may signal how they expect to play, and thereby alleviate the tension due to strategic uncertainty.[2] The auctions also capture important aspects of "general equilibrium" analogs of VHBB's earlier environments, in which players choose among coordination games with the market-clearing price analogous to the opportunity costs determined by their best alternatives.

Auctioning the right to play the median game had remarkable consequences. When that game was played without auctions in VHBB's (1991) experiment, most subjects initially chose inefficiently low efforts, and six out of six subject groups converged to inefficient equilibria. Auctions might be expected to yield more efficient outcomes simply because subjects have diverse beliefs about each other's efforts, auctions select the most optimistic subjects to play, and optimism favors efficiency in the median game. But VHBB's (1993) subjects did much better than this argument suggests: In eight out of eight groups, they bid the market-clearing price to a level recoverable only in the efficient equilibrium and then converged to that equilibrium.[3] Convergence was very fast, essentially complete within three to five periods. The limiting outcome was consistent with subgame-perfect equilibrium in the *stage game* consisting of the auction followed by the coordination game, and subjects' behavior suggests that their beliefs were focused as in the intuition for forward induction refinements, in which players infer from their partners' willingness to pay to play a game that they expect their payoffs to repay their costs, and intend to play accordingly (Elchanen Ben-Porath and Eddie Dekel, 1992).[4]

The efficiency-enhancing effect of auctioning the right to play in VHBB's (1993) experiment suggests a novel and potentially important way in which competition might promote efficiency. Although it conveys a powerful impression by itself, it raises as many questions as it answers. Was the outcome

[2] The auction is an unusual form of preplay communication in that players' messages can directly influence their payoffs, hence are not "cheap talk"; and they are communicated only through an aggregate, the market-clearing price.

[3] Roberto Weber (1994) has replicated this result in a closely related environment, and Gerard Cachon and Colin Camerer (1996) have verified its robustness and refined its interpretation. Their experimental designs and results, and the extension of our analysis to explain them, are discussed in Crawford and Broseta (1995).

[4] The intuition for forward induction seems to favor the efficient equilibrium, and subjects' limiting behavior was consistent with forward induction, as it is usually formalized, as well as subgame perfectness. But both refinements are consistent with any of the seven symmetric pure-strategy equilibria in the coordination game (with full surplus extraction in the auction), so neither helps to explain VHBB's result.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

VHBB observed inevitable in their environment? How would the strength of the effect vary with the *treatment variables*—the number of players in the auction and the coordination game and the order statistic that determines the robustness of the efficient equilibrium—which never varied in the experiment? Would the effect extend to environments beyond those that directly generalize VHBB's designs?

In principle such questions could be answered piecemeal by further experiments, but only a theory that elucidates the mechanism behind VHBB's result can provide a firm basis for generalization and realize its full power to inform analysis. Such a theory can change the way we think about many applications in which the participants in coordination games are determined by a sorting process like VHBB's auction or its general equilibrium analogs.

The models in Cooper and John (1988), for instance, view the entire economy as a coordination game, but it may be more realistic to view it as composed of sectors, regions, or firms, each of which is a coordination game. Because participants must often choose among these, the economy may be closer to VHBB's auction environment than the games with exogenous participation of their earlier experiments. If so, VHBB's result suggests that modeling the economy as a single coordination game systematically overstates the power of coordination failure to explain underemployment. Alternatively, consider a coordination model of Hong Kong's economy that takes into account the high costs most inhabitants paid to get there.[5] Did Hong Kong do so well by attracting entrants who expected efficient coordination, as in our "optimistic subjects" effect; via the power of entry barriers to focus beliefs on better coordination outcomes, as in our "forward induction" effect; by having rules that favor efficient coordination, perhaps as in the "robustness" effect discussed below; or simply by exploiting its natural advantages and "skimming the cream" from mainland China? Only a theory that distinguishes and quantifies these effects can give convincing answers to such questions.

In this paper we propose a theory to explain VHBB's (1993) result, generalizing our analyses of VHBB's (1990, 1991) results in Crawford (1995) and Broseta (1993a, 1993b) to a class of coordination games with auctions in which VHBB's treatment variables can take arbitrary values.[6] The main difficulty our analysis must address is explaining equilibrium selection; the rapid convergence to equilibrium in the stage game in VHBB's (1993) experiment and most of their earlier treatments is easily explained by learning models.

In VHBB's (1990, 1991) experiments, the dynamics and limiting outcomes varied systematically with the size of the groups playing the game and the order statistic that determined their payoffs and best responses, with strong downward drift and convergence to the equilibrium with lowest effort in VHBB's (1990) large-group minimum games, but no drift and consistent "lock-in" on the initial median in VHBB's (1991) median games. These and other variations in equilibrium selection across VHBB's treatments discriminate sharply among alternative theories of strategic behavior, both traditional and adaptive.

Our analyses explained these variations as persistent effects of interactions between strategic uncertainty and history-dependent learning dynamics. One would expect the level of strategic uncertainty to decline gradually to zero as players learn to forecast each other's responses. Unless this decline is very slow, the learning dynamics lock in on a particular equilibrium in the limit, and the model's implications for equilibrium selection can be summarized by the prior probability distribution of that equilibrium, which is normally nondegenerate due to the persistent effects of strategic uncertainty. The distribution of the

[5] We are grateful to a co-editor for suggesting this illustration, which is closer to VHBB's auction environment because Hong Kong's position has been nearly unique.

[6] Yong-Gwan Kim (1996) suggests interesting alternative explanations of VHBB's result, based on adaptive dynamics and axioms in the spirit of evolutionary stability.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

limiting equilibrium is determined by players' learning rules, the treatment variables, and the cumulative effects of strategic uncertainty on the dynamics. The key to our explanation is that strategic uncertainty generates what we shall call a robustness effect, which gives the learning dynamics a negative, zero, or positive drift for order statistics respectively below, at, or above the median. The magnitude of the drift increases with the level of strategic uncertainty, and for sufficiently large initial levels, declining over time, it makes the prior probability distribution of the limiting equilibrium vary across VHBB's (1990, 1991) treatments much as its empirical frequency distribution varied in their experiments.

Here we adapt our earlier methods to show that in VHBB's (1993) environment, strategic uncertainty generates robustness, optimistic subjects, and forward induction effects large enough to explain the pattern of equilibrium selection they observed. Our approach yields a unified explanation of VHBB's (1990, 1991, 1993) results, and makes it possible to assess the likely importance of those effects in other environments.

We assume throughout that players are rational in the standard sense that their decisions maximize expected payoffs, given their beliefs about each other's decisions. But our analysis is adaptive in that, rather than assuming equilibrium in the stage game or the repeated game that describes players' entire interaction, we allow players to have diverse beliefs about each other's decisions in the stage game and model the process by which their beliefs and decisions converge as they learn to predict each other's decisions. Thus, our analysis allows equilibrium to be reached by unsophisticated learning from actual observation rather than introspective, strategically sophisticated reasoning.

We impose the structure needed for analysis by using special features of VHBB's environment to give a simple econometric characterization of strategic uncertainty and the learning dynamics, which allows a more informative analysis of the effects of strategic uncertainty than now seems possible for games in general. Players' initial beliefs and responses to new observations each period are perturbed by independently and identically distributed (henceforth, "i.i.d.") random shocks. The shocks represent strategic uncertainty in terms of the differences in players' learning rules, and their variances represent the initial level of strategic uncertainty and how it varies over time as players learn to forecast each other's responses. In effect each player has his own theory of coordination, which gives his initial beliefs and his interpretations of new observations an unpredictable component.

The resulting model is a Markov process with time-varying transition probabilities, the dynamics of which are driven by strategic uncertainty. Like our earlier models, it encompasses the leading alternative characterizations of strategic behavior in such settings: traditional analyses in which players' beliefs and decisions are focused from the start on a particular equilibrium, evolutionary analyses of the "long-run equilibria" of ergodic dynamics with small amounts of noise, and history-dependent learning with lock-in on a particular equilibrium in the limit. These are distinguished mainly by different values of the variances that represent the level of strategic uncertainty and how it varies over time.

As in our earlier analyses, we take the variances and certain other aspects of behavior to be exogenous behavioral parameters. Our characterization provides a framework within which to close the model by estimating them, using VHBB's (1993) experimental data.[7] As before, the estimated parameters generally satisfy the restrictions suggested by theory, but differ significantly from the values needed to justify an equilibrium analysis or an analysis of the long-run equilibria of ergodic dynamics. Instead they indicate large initial levels of strategic uncertainty, declining to zero over time. The decline is rapid enough to make the

[7] The need to proceed this way will come as no surprise to anyone who accepts Thomas Schelling's (1960) premise that the analysis of coordination is inherently partly empirical. Strategic uncertainty is crucial to understanding the dynamics, but the differences in subjects' beliefs cannot be reliably explained by theory alone because subjects were indistinguishable and had nearly identical information.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

dynamics converge with probability one to one of the pure-strategy subgame-perfect equilibria of the stage game, so that the model's implications for equilibrium selection can be summarized by the prior probability distribution of that equilibrium. The estimated model gives an adequate statistical summary of individual subjects' behavior, and implies probability distributions of the dynamics and limiting outcome that resemble their empirical frequency distributions in VHBB's experiment.

We study equilibrium selection in more detail by using special features of our learning model and the environment to obtain a closed-form solution for the histories of players' beliefs, bids, and efforts as functions of the behavioral parameters, the treatment variables, and the shocks that represent strategic uncertainty. The solution shows how the limiting outcome is built up period by period from the shocks, the effects of which persist indefinitely. This persistence makes the learning process resemble a random walk in the aggregate, but with declining variances and nonzero drift that depend on the parameters and treatment variables. As in our earlier analyses, large, persistent effects of strategic uncertainty make it necessary to analyze the entire learning process to understand equilibrium selection, and preclude explanations based on equilibrium refinements or ergodic dynamics, which assume away either strategic uncertainty or persistence.

The form of our solution indicates that unless the behavioral parameters vary sharply and unpredictably with changes in the environment—which our estimates for VHBB's (1990, 1991) experiments suggest is unlikely—the dynamics and limiting outcome will vary with the treatment variables in stable, predictable ways. We begin with a qualitative comparative dynamics analysis of the effects of changes in VHBB's (1993) treatment variables. Generalizing results from our earlier analyses, we make precise, in the probabilistic sense appropriate to the model, the common intuitions that coordination tends to be less efficient the closer the order statistic is to the minimum (our robustness effect) and less efficient in larger groups because it requires coherence among a larger number of independent decisions. We also establish a new result, showing that coordination tends to be more efficient the more intense the competition for the right to play.

Our model explains the efficiency-enhancing effect of auctions as the result of a dynamic interplay between forward induction, optimistic subjects, and robustness effects. Our solution allows a quantitative comparative dynamics analysis that shows how the magnitudes of these effects are determined by the treatment variables and behavioral parameters, and makes it possible to estimate their importance in other environments. The mean coordination outcome can be expressed as the sum of four components per period, one of which is the forward induction effect, one of which combines the optimistic subjects effect with the robustness effect of our earlier analyses, and two of which are smaller effects, discussed below. These components can be approximated as known functions of the treatment variables, the behavioral parameters, and tabulated statistical parameters.

These approximations yield a simple characterization of the optimistic subjects and robustness effects. Together they have approximately the same magnitude in VHBB's environment (where the right to play a nine-person median game was auctioned in a group of 18) as the robustness effect in an 18-person coordination game without auctions in which payoffs and best responses are determined by the fifth highest (the median of the nine highest) of all 18 players' efforts. In this respect the auctions effectively transformed VHBB's median game, which without auctions would have a robustness effect that contributes zero drift to the dynamics, into a 75th-percentile game ($0.75 = 13.5/18$) with a robustness effect that contributes a large upward drift. Our estimates suggest that this drift is responsible for roughly half of the efficiency-enhancing effect of auctions in VHBB's environment, and that the other half is due to a strong forward induction effect.

Our quantitative analysis also shows that unless the behavioral parameters vary sharply with changes in the environment, there is a positive efficiency-enhancing effect through-

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

out the class of environments we study, but it is not always strong enough to assure fully efficient coordination. More generally, the mechanism suggests that the effect will extend to other laboratory or field environments that combine significant strategic uncertainty with the ingredients of our optimistic subjects and forward induction effects.

The rest of the paper is organized as follows. Section I describes the environments we study and introduces our learning model. Sections II and III present the theoretical and econometric analyses. Section IV is the conclusion.

I. The Model

In this section we describe the class of environments we study and introduce the learning model that is the basis for our theoretical and econometric analyses. We call the highest price at which a player plans to remain in the auction his *bid,* and we define *efficiency* with reference to players' payoffs in the coordination game, bearing in mind that the auction may transfer some or all of the surplus to the experimenters.

A. *The Environment*

In VHBB's (1993) experiment subjects played the nine-person median coordination game of VHBB's (1991) treatment Γ for 10 or 15 periods, with the right to play auctioned each period in a population of 18 subjects. In the median game nine players choose among seven efforts, $\{1, \dots, 7\}$, with payoffs determined by their own efforts and the group median effort. Denoting players' efforts at time t by $\hat{x}_{1t}, \dots, \hat{x}_{nt}$ and the median by M_t, player i's period-t payoff in dollars is $0.1M_t - 0.05(M_t - \hat{x}_{it})^2 + 0.6$. Because no player's effort can influence the median when all other players are choosing the same effort, and player i's payoff is higher, other things equal, when $\hat{x}_{it} = M_t$, any configuration of efforts with $\hat{x}_{it} = M_t$ for all i is an equilibrium, and these seven symmetric equilibria are the only pure-strategy equilibria. These equilibria are strict and Pareto ranked, with equilibria with higher M_t better for all players than those with lower M_t.

The auction was a multiple-unit ascending-bid English clock as in Kevin McCabe et al. (1990), in which subjects indicated their willingness to pay the current asking price by holding up bid cards, with subjects who dropped out not allowed to reenter the bidding. The asking price started five cents below the payoff of the least efficient equilibrium of the median game and increased by five-cent increments until 11 or fewer subjects remained, after which it increased by one-cent increments until nine or fewer subjects remained. (The payoffs of adjacent equilibria differed by ten cents.) The market-clearing price was determined as follows. If the lowest price at which nine or fewer subjects remained in the auction left exactly nine subjects, all nine were awarded the right to play at that price. If that price left fewer than nine subjects, they were all awarded the right to play, with the remainder of the nine slots filled randomly from those who dropped out at the last increase, and all nine subjects paying the price before the last increase.[8]

In VHBB's (1993) design, as in their treatment Γ, explicit communication was prohibited throughout; the median was publicly announced after each play; and with minor exceptions the structure was publicly announced at the start. The market-clearing price was publicly announced after each auction, before the winners played the median game.

We now describe a class of environments that generalizes VHBB's (1993) design by allowing any values of their treatment variables and any coordination game that shares the structural features noted above (including those used in VHBB's [1990, 1991] other treatments). We assume throughout that bids are continuously variable.[9] We introduce the model under the simplifying assumption that effort is also continuously variable. In Section I, subsection D, and more formally in Section III, subsection B, we explain how to adapt the

[8] Thus subjects never paid more than they indicated they were willing to, and they were never excluded involuntarily unless they had indicated approximate indifference.

[9] This was not literally true in the experiments, but the increments by which bids could vary were small enough (one cent near the market-clearing price, compared to a ten-cent payoff difference between equilibria) to make it a good approximation.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

model to the discrete effort spaces of VHBB's experiment by viewing the continuously variable efforts as the latent variables in an ordered probit model of discrete effort choice.

There is a finite number, m, of indistinguishable players, who repeatedly play a stage game with symmetric player roles. The stage game consists of an n-person coordination game, with $n < m$, preceded by an auction in which all m players bid for the right to play.

The coordination game has one-dimensional strategy spaces, with strategies called efforts. Effort has a commonly understood scale, which makes it meaningful to say that a player chose the same effort in different periods, or that different players chose the same effort. Any symmetric effort combination is an equilibrium, and such combinations are the only pure-strategy equilibria. By symmetry, these equilibria are Pareto ranked unless players are indifferent between them. Each player's best responses are given by a summary statistic of all players' efforts when that statistic is unaffected by his effort. The summary statistic, y_t, is a function $f(x_{1t}, \dots, x_{nt})$, where x_{it} is player i's continuously variable effort at time t and players $1, \dots, n$ play the game. We assume that $f(\cdot)$ is continuous and, for any $x_{1t}, \dots, x_{nt}$ and constants a and $b \geq 0$, $f(a + bx_{1t}, \dots, a + bx_{nt}) \equiv a + bf(x_{1t}, \dots, x_{nt})$. These assumptions are satisfied when $f(\cdot)$ is an order statistic or a convex combination of order statistics such as the arithmetic mean. To see what they entail, note that the symmetry of roles implies that $f(\cdot)$ is a symmetric function of the x_{it}, so that its value is determined by the order statistics of their empirical distribution. Our assumptions rule out most nonlinear functions of these order statistics, which is restrictive but probably not unrepresentative of symmetric games. We abuse terminology by calling $f(\cdot)$ an order statistic.

The auction is a multiple-unit ascending-bid English clock as described above, with player i's bid and the market-clearing price at time t denoted p_{it} and q_t, respectively. We depart from the auction VHBB used only in assuming that p_{it} and q_t are continuously variable. With continuously variable bids the effects of VHBB's tie-breaking procedure are negligible, and q_t can be taken to be the $(n + 1)$th largest of the p_{it}, an order statistic of their empirical distribution. For any given m and n, q_t can be written as a function,

$$(1) \qquad q_t \equiv g(p_{1t}, \dots, p_{mt}),$$

where for any $p_{1t}, \dots, p_{mt}$ and any constants a and $b \geq 0$,

$$(2) \qquad g(a + bp_{1t}, \dots, a + bp_{mt}) \equiv a + bg(p_{1t}, \dots, p_{mt}).$$

Because the form of $g(\cdot)$ is completely determined by m and n, we describe changes in $g(\cdot)$ below by describing the associated changes in m and n, without separate reference to $g(\cdot)$.

The outcome of the stage game can now be described more precisely. We write

$$(3) \qquad y_t \equiv h(x_{1t}, \dots, x_{mt}; p_{1t}, \dots, p_{mt}),$$

where the x_{it} now include planned effort choices for all m players, and $h(x_{1t}, \dots, x_{mt}; p_{1t}, \dots, p_{mt})$ equals $f(\cdot)$ evaluated at the x_{it} associated with the n highest p_{it} in the population—noting that the order of these x_{it} is immaterial and ignoring ties in the p_{it}, which will have probability zero under our distributional assumptions. It is immediate from this definition that for any $p_{1t}, \dots, p_{mt}$ and any constants a and $b \geq 0$, $h(\cdot)$ is continuous in $x_{1t}, \dots, x_{mt}$ and inherits the scaling properties of $f(\cdot)$, in that

$$(4) \qquad h(a + bx_{1t}, \dots, a + bx_{mt}; p_{1t}, \dots, p_{mt}) \equiv a + bh(x_{1t}, \dots, x_{mt}; p_{1t}, \dots, p_{mt}).$$

As only the order of the p_{it} matters, it is also clear that for any constants a and $b > 0$,

$$(5) \qquad h(x_{1t}, \dots, x_{mt}; a + bp_{1t}, \dots, a + bp_{mt}) \equiv h(x_{1t}, \dots, x_{mt}; p_{1t}, \dots, p_{mt}).$$

We assume that the structure of the environment is publicly known.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

B. *Equilibrium Refinements in the Stage Game*

As noted above, in VHBB's (1993) experiment eight out of eight subject groups bid the market-clearing price of the right to play to a level recoverable only in the efficient equilibrium, and their efforts converged to that equilibrium. As VHBB's (1993) Tables V and VI make clear, subjects' bids and efforts were generally consistent with the intuition for forward induction, in that individual subjects seldom bid more than their efforts made it possible to recover in the coordination game, and never did so after the first few periods.

Although we shall argue below that the limiting outcomes in the experiment cannot be understood without analyzing history-dependent learning dynamics, it is instructive to consider the implications of equilibrium refinements in the stage game. Because the structure was publicly announced, we assume complete information. The most relevant refinements are subgame perfectness and forward induction, and the essential points can be seen by focusing on symmetric pure-strategy equilibria. Without loss of generality we rescale the payoffs in the coordination game, which are increasing in players' common efforts, so that in each of these equilibria players' payoffs equal their efforts.

When all players use the same bidding strategy and $n < m$, no player can alter the market-clearing price by changing his bid. It is easy to check that any strategy combination with $p_{it} = x_{it} = q_t = y_t$ for all i, in which each player bids his payoff in a symmetric pure-strategy equilibrium of the coordination game and then plays that equilibrium's effort unless he is the only player to bid that amount, is a subgame-perfect equilibrium of the stage game.[10]

In applying forward induction we follow Ben-Porath and Dekel (1992) in identifying it with the ability to survive iterated deletion of weakly dominated strategies. VHBB's stage game resembles Ben-Porath and Dekel's Figure 2.3b example, in which two players publicly and simultaneously decide whether to burn a given quantity of money before playing a coordination game like VHBB's.[11] This game has a symmetric subgame-perfect equilibrium in which players burn money and then play the inferior equilibrium in the coordination game, each anticipating that equilibrium unless he is the only player who burns money. In this equilibrium each player's strategy is a unique best response given his beliefs, so it survives iterated deletion of dominated strategies. As Ben-Porath and Dekel note, such equilibria normally exist whenever players move simultaneously in the communication phase. In VHBB's stage game such an equilibrium can be constructed to support any of the seven symmetric pure-strategy equilibria in their coordination game, with full surplus extraction in the auction. The requirements for such equilibria are also consistent with subgame perfectness. Thus, subgame perfectness and forward induction in the stage game, together or separately, are consistent with the limiting outcomes in VHBB's experiment but too unrestrictive to help in explaining them. We use "forward induction" loosely below to refer to the intuition rather than a formal definition.

C. *The Learning Model*

For our purposes, it is essential to capture the idea that, even if players form their beliefs and choose their bids and efforts sensibly, they may differ in unpredictable ways. It is also essential to describe the dynamics of beliefs,

[10] No player can gain by changing his bid because the right to play yields zero surplus. Other symmetric pure-strategy subgame-perfect equilibria differ from these only inessentially, off the equilibrium path. There are no symmetric equilibria in which the right to play yields positive surplus because a player could gain by outbidding his partners, assuring himself the right to play without raising the market-clearing price. Kim (1996 Lemma 1) gives a full characterization of the pure-strategy subgame-perfect equilibria in this game, including asymmetric ones.

[11] VHBB's stage game differs from Ben-Porath and Dekel's example in that the latter's players must play the game whether or not they burn money, and must bear the cost of any money they burn no matter what their partners do; but these differences are inessential here. Note that unilateral deviations from symmetric bid combinations are always observable, although deviations from asymmetric combinations might not be.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

bids, and efforts realistically in terms of observable variables, in a way that permits estimation of the behavioral parameters that cannot reliably be determined by theory (such as the variances that represent the level of strategic uncertainty) and allows an informative theoretical analysis.

In the environments we study, players' bids and efforts evolve as follows. First all m players choose their initial bids, the resulting value of q_t is publicly announced, and the n winners choose their efforts in the coordination game.[12] All m players then observe the resulting value of y_t and choose new bids, and the process continues. Because the structure is public knowledge, players face uncertainty only about each other's bids and efforts. The effects of others' bids and efforts on a player's payoffs are filtered through q_t and y_t.

The large subject populations in VHBB's (1993) experiment make it a plausible working hypothesis that players treat their individual influences on q_t and y_t as negligible. This implies that players' optimal bids and efforts each period are determined by their current payoff implications, and thus by their beliefs about the current y_t, because those beliefs determine their expected payoffs and optimal efforts, and their bids make winning the right to play contingent on q_t. We imagine that players begin with prior beliefs about the process that generates q_t and y_t, use standard statistical procedures to revise their beliefs in response to new observations, and choose the bids and efforts that are optimal, given their beliefs. Players whose priors differ may then have different beliefs even if they have always observed the same history and used the same procedures to interpret it.

We adopt a simple econometric characterization of the dynamics of beliefs, bids, and efforts in the style of the adaptive control literature.[13] The key insight of the control literature is that describing how beliefs respond to new information does not require representing them as probability distributions or their moments. It is enough to model the dynamics of the optimal decisions they imply, which are the only aspect of beliefs that directly affects the outcome. We represent players' beliefs by their optimal efforts, which when continuously variable preserve enough information to realistically describe their bids as well as their efforts.[14] There is no need to assume that the optimal efforts are determined in any simple way by the moments of the distributions that describe players' beliefs.

We describe the conditional means of players' responses to new observations by simple linear adjustment rules, in which their beliefs adjust part of the way toward the value suggested by the latest observation of y_t in a way that generalizes the fictitious-play and best-response rules to allow different values of parameters that represent the initial levels, trends, and inertia of beliefs. We describe the differences in players' beliefs by perturbing these mean adjustments by idiosyncratic random shocks each period, which represent strategic uncertainty described in terms of the differences in players' learning rules. We assume that these shocks are i.i.d. across players, with zero means and given variances. Thus, any correlation in players' beliefs that emerges is the result of responses to common observations of y_t and q_t rather than an artifact of our distributional assumptions.

We distinguish between players' beliefs when they bid and when the winners in the auction play the coordination game, because the latter beliefs may reflect inferences from q_t. Player i's beliefs in period t before and after he observes q_t are denoted $\bar{x}_{it}$ and x_{it}, respectively. They are interpreted as his estimates of his optimal effort given his information at those times, with $\bar{\mathbf{x}}_t$ and $\mathbf{x}_t$ denoting the associated population m-vectors.

[12] We assume that bids are not publicly observable, even though VHBB's subjects were in booths that may have allowed them, with some effort, to observe when other subjects dropped out. At our request Weber (1994) ran two trials that replicated VHBB's design but ensured that bids were unobservable, with very similar results.

[13] See M. B. Nevel'son and R. Z. Has'minskii (1973) or Michael Woodford (1990 Sec. 2).

[14] Our model with discrete efforts, by contrast, describes beliefs and efforts separately.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Players' initial beliefs are given by

$$(6)\qquad \bar{x}_{i0} = \alpha_0 + \zeta_{i0},$$

where α_0 is their common mean and the ζ_{i0} are idiosyncratic shocks.

Recalling that payoffs in the coordination game are measured so that players' payoffs in symmetric equilibria equal their efforts, players' bids in period t are given by

$$(7)\qquad p_{it} = \beta_t + \bar{x}_{it} + \eta_{it}, \qquad t = 0, \dots,$$

where the β_t are constants and the η_{it} are idiosyncratic shocks that represent the differences in players' bidding strategies. This specification can be justified as follows. Players' values of the right to play may be affiliated, because with strategic uncertainty higher bids signal higher beliefs and efforts, which raise the *ex ante* distributions of all players' payoffs.[15] But because players ignore their future influences on y_t, and their individual influences on q_t are negligible unless they do not win the right to play, standard arguments show that it is approximately optimal for a player to stay in the auction until his value given the current level of q_t falls below q_t [although q_t influences the outcome in the game, via (8) below, this criterion uniquely defines his bid when $\delta_t < 1$].[16] Finally, because VHBB's median game imposes a quadratic payoff penalty for efforts away from the median, a risk-neutral player's optimal effort equals the mean of his distribution of y_t by certainty-equivalence. His expected payoff is then equal to that effort minus the variance of his distribution of y_t, which should be approximately proportional to the variance of θ_{it} in (8) below. This yields (7), with β_t, the mean of players' corrections for the cost of strategic uncertainty, negative and approaching zero over time if Var $\theta_{it} \rightarrow 0$.

After the auction players observe q_t, as determined by (1), and update their beliefs to

$$(8)\qquad x_{it} = \gamma_t + \delta_t q_t + (1 - \delta_t)\bar{x}_{it} + \theta_{it}, \qquad t = 0, \dots,$$

where the γ_t and δ_t are constants and the θ_{it} are shocks that represent the differences in updating rules. γ_t and δ_t, the level and slope parameters of these linear adjustments, together reflect forward induction inferences from the observed value of q_t.

The winners then determine their efforts according to the x_{it}, and all m players observe the value of y_t determined by (3) and update their beliefs again, to

$$(9)\qquad \bar{x}_{it+1} = \alpha_{t+1} + \varepsilon_{t+1} y_t + (1 - \varepsilon_{t+1}) x_{it} + \zeta_{it+1}, \qquad t = 0, \dots,$$

where the α_t and ε_t are constants and the ζ_{it} are shocks. The level and slope parameters α_{t+1} and ε_{t+1} represent any trends in beliefs and the precision of beliefs, as reflected by how much they respond to new observations of y_t. The $\bar{x}_{it+1}$ then determine players' bids next period via (7), and the process continues.

Although (8) and (9) suggest partial adjustment to the efforts suggested by the latest observation of q_t or y_t, they are best thought of as representing full adjustment to players' current estimates of their optimal efforts, which respond less than fully to new observations because they are only part of players' information about the process. As explained in Crawford (1995 pp. 112–13), (9) generalizes the familiar fictitious-play and best-

[15] Paul Milgrom and Robert Weber (1982) define affiliation. Even though winners play the same game, their values are not common because their beliefs and efforts differ.

[16] VHBB's auction is a multiple-unit version of the "Japanese" auction studied by Milgrom and Weber (1982). They show that when players have affiliated values and can observe each other's bids, it is optimal for a player to stay in until the current price makes him indifferent between winning or losing, taking into account the winner's curse. However, in VHBB's auctions winning reveals nothing to a player about other winners' beliefs; and anything winners learn about losers' beliefs is irrelevant to the current value of the game unless the idiosyncratic parts of players' beliefs are correlated. This suggests that winners' curse effects are unimportant here, and that it makes little difference whether players can observe losers' bids.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

response learning rules to allow a much wider range of priors about the structure of the y_t process.[17] Equations (8) and (9) are not fully "rational" in the game-theoretic sense, because players' priors are not required to be correct and linearity may be inconsistent with the rules that are optimal given their priors. However, in the control literature such rules have been shown to provide a robust estimation method for agents who understand the forecasting problems they face, but are unwilling to make the specific assumptions about the structure of the process or unable to store and process the large amounts of information a Bayesian approach requires.

By construction, $E\zeta_{it} = E\eta_{it} = E\theta_{it} = 0$ for all i and t, where E is the expectation operator. We assume that ζ_{it}, η_{it}, and θ_{it} are independent of each other and serially independent for each i, which amounts to assuming that $\bar{x}_{it}$ and x_{it} fully capture the future effects of past idiosyncratic influences on player i's beliefs. We also assume that ζ_{it}, η_{it}, and θ_{it} are i.i.d. across i, with common variances denoted $\sigma^2_{\zeta t}$, $\sigma^2_{\eta t}$, and $\sigma^2_{\theta t}$.

We assume throughout that $0 < \delta_t, \varepsilon_t \leq 1$. With no further restrictions on the behavioral parameters, the model is formally consistent with any history of the y_t for any m, n, and $f(\cdot)$, with α_t varying over time as necessary and β_t, γ_t, $\sigma^2_{\zeta t}$, $\sigma^2_{\eta t}$, and $\sigma^2_{\theta t}$ set equal to zero, so that for all i and t, $p_{it} = \bar{x}_{it} = x_{it} = q_t = y_t = \Sigma^t_{s=0}\, \alpha_s$. Such solutions, in which players jump from one stage-game equilibrium to the next following some commonly understood pattern, are not ruled out by standard arguments because they correspond to equilibria of the repeated game. Yet they are empirically bizarre, because the ad hoc time variation in the behavioral parameters they require violates the hypothesis that, in a given environment, the distribution of past behavior is indicative of current behavior, which most of us take for granted in learning to predict others' behavior in repeated interactions.

We address this issue below by imposing intertemporal restrictions that rule out such ad hoc variations in the model's behavioral parameters. These restrictions allow for the possibility that α_t, β_t, γ_t, $\sigma^2_{\zeta t}$, $\sigma^2_{\eta t}$, $\sigma^2_{\theta t} \to 0$ as $t \to \infty$, so that players learn to predict y_t if it converges and choose the x_{it} that are optimal given their predictions, and the long-run steady states of the dynamics coincide with the pure-strategy subgame-perfect equilibria of the stage game [shown in Section I, subsection B, to satisfy $p_{it} = x_{it}$ $(= \bar{x}_{it}) = q_t = y_t$ for all i]. We find below that they permit an adequate econometric description of individual subjects' behavior, which yields a useful explanation of the dynamics of their interactions.

D. *Discrete Efforts*

To estimate the model and to use it to predict experimental results, we must allow for the discrete efforts of VHBB's design. We do this by letting x_{it} and $\bar{x}_{it}$ continue to represent beliefs and remain continuously variable, but with x_{it} determining effort as the latent variable in an ordered probit, described in Section III, subsection B. Beliefs, bids, and efforts still follow (1), (3), and (6)–(9), but with efforts and the order statistic determined by rounding each x_{it} to the nearest feasible integer and evaluating $h(\cdot)$ at the rounded values.

II. Analysis

We begin by analyzing the model with continuously variable efforts. In Section II, subsection F, we explain how to extend the analysis to the discrete efforts of VHBB's experiment.

A. *Preliminaries*

Once the distributions of ζ_{it}, η_{it}, and θ_{it} are specified, (1), (3), and (6)–(9) define a Markov process with state vector $\bar{\mathbf{x}}_t$, in which beliefs, bids, and efforts are i.i.d. across players *ex ante*. Equation (6) provides initial conditions and substituting (1),

[17] Unlike learning rules in which players respond only to realized payoffs, such as those in Alvin Roth and Ido Erev (1995), (8) and (9) reflect the best-response structure. As explained in Crawford (1997 Sec. 6.3), VHBB's subjects seemed to understand the structure, and such rules allow a better description of their behavior.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

(3), and (7)–(8) into (9) and using (2) and (4)–(5) gives $\bar{x}_{it+1}$ as a function of the $\bar{x}_{it}$, parameters, and shocks:

$$
\begin{aligned}
(11)\quad \bar{x}_{it+1} = {} & \alpha_{t+1} + \beta_t\delta_t + \gamma_t \\
& + \delta_t g(\bar{x}_{1t} + \eta_{1t}, \dots, \bar{x}_{mt} + \eta_{mt}) \\
& + \varepsilon_{t+1} h[(1-\delta_t)\bar{x}_{1t} + \theta_{1t}, \dots, \\
& (1-\delta_t)\bar{x}_{mt} + \theta_{mt}; \bar{x}_{1t} + \eta_{1t}, \dots, \bar{x}_{mt} + \eta_{mt}] \\
& + (1-\varepsilon_{t+1})[(1-\delta_t)\bar{x}_{it} + \theta_{it}] \\
& + \zeta_{it+1}, \qquad t = 0, \dots .
\end{aligned}
$$

The model's recursive structure and the independence of players' deviations from the average adjustment rules capture the requirement that players form their beliefs and choose their strategies independently that is the essence of the coordination problem.

The model's dynamics are driven by strategic uncertainty, as represented by the $\sigma^2_{\zeta t}$, $\sigma^2_{\eta t}$, and $\sigma^2_{\theta t}$. Different assumptions about how $\sigma^2_{\zeta t}$, $\sigma^2_{\eta t}$, and $\sigma^2_{\theta t}$ vary over time have different implications for the dynamics, which go a long way toward identifying the stochastic structure. The following result is helpful in understanding this relationship.

PROPOSITION 0: *Suppose that $\alpha_t = 0$ for all $t = 1, \dots$ and $\beta_t = \gamma_t = 0$ for all $t = 0, \dots$, so that there are no exogenous trends in beliefs, and that $\sigma^2_{\eta t} = \sigma^2_{\theta t} = \sigma^2_{\zeta t} = 0$ for all $t = T, \dots$, so that the rules that describe players' responses to new observations coincide from period T onward for some $T \geq 0$. Then there exist players, j and k, such that $g(\bar{x}_{1t}, \dots, \bar{x}_{mt}) = \bar{x}_{jt}$ and $h(\bar{x}_{1t}, \dots, \bar{x}_{mt}; \bar{x}_{1t}, \dots, \bar{x}_{mt}) = \bar{x}_{kt}$ for all $t = T, \dots$; and $\bar{x}_{it}$ and x_{it} converge with probability one to a common value, weakly between $\bar{x}_{jT}$ and $\bar{x}_{kT}$, that is completely determined by $\bar{x}_{jT}$, $\bar{x}_{kT}$, and the δ_t and ε_t for $t = T, \dots$. If $T = 0$, so that players' initial beliefs coincide as well, then $\bar{x}_{it} = x_{it} = \alpha_0$ for all i and all $t = 0, \dots$.*

The proof is in the Appendix. Proposition 0 shows that if there are no differences in players' responses to new information from period T on, the dynamics must converge to an equilibrium dictated by, and bounded between, the $\bar{x}_{jT}$ and $\bar{x}_{kT}$ that determine $g(\cdot)$ and $h(\cdot)$ in period T. In VHBB's environment $\bar{x}_{jT}$ and $\bar{x}_{kT}$ are the tenth and fifth largest $\bar{x}_{iT}$ (because $n = 9$, $g(\cdot)$ is the $(n+1)$th highest bid, and $h(\cdot)$ is the median of the nine largest $\bar{x}_{iT}$). Thus, unless differences in responses persist long enough to drive $\bar{x}_{jT}$ and $\bar{x}_{kT}$ to efficient levels, the model cannot explain the convergence to the efficient equilibrium VHBB observed.[18]

If $\sigma^2_{\eta t}$, $\sigma^2_{\theta t}$, and $\sigma^2_{\zeta t}$ converge to positive constants, instead of remaining at 0 for all $t = T, \dots$, an analysis of long-run equilibria seems possible, along the lines of the analysis of evolutionary dynamics in Kim (1996). This would probably reproduce Kim's conclusion that there is efficient coordination in the long run, as in VHBB's experiment. However, the closely related experiments of Weber (1994) and Cachon and Camerer (1996) suggest that efficiency is not inevitable for all values of m, n, and $f(\cdot)$ (although no amount of data from experiments with finite horizons can ever disprove a long-run result). More importantly, we find below that when $\sigma^2_{\eta t}$, $\sigma^2_{\theta t}$, and $\sigma^2_{\zeta t}$ decline gradually to 0, as one would expect as players learn to forecast q_t and y_t from their common observations, the model yields history-dependent dynamics that closely resemble VHBB's results. And in this case the model also provides useful information about the mechanism behind the efficiency-enhancing effect of auctions and the effects of changes in the environment—information that would very likely be lost in an analysis of long-run equilibria.

B. *Closed-Form Solution*

We now show how to analyze the dynamics, whether or not the variances converge to zero, in a way that yields a deeper understanding of

[18] This not the whole story because in our empirical analysis α_t, β_t, and γ_t sometimes differ from 0 and, as explained below, these parameters also affect efficiency. We set them equal to 0 here, except for α_0, to focus on a more subtle aspect of the dynamics.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

how outcomes are determined and makes it possible to estimate the effects of changes in the environment.

The key to much of our analysis is the fact that the scaling properties of the order statistics $g(\cdot)$ and $h(\cdot)$ and the form of players' learning rules allow us, despite the model's nonlinearity, to obtain a closed-form solution for the entire history of players' beliefs, bids, and efforts as a function of the behavioral parameters and the shocks that represent strategic uncertainty. In what follows, sums with no terms (like $\sum_{s=0}^{t-1} \varepsilon_{s+1} h_s$ for $t = 0$) are understood to equal 0, and products with no terms are understood to equal 1.

PROPOSITION 1: *The unique solution of the learning dynamics is given, for all i and t, by*

$$(12) \quad x_{it} = \sum_{s=0}^{t} [\alpha_s + \beta_s\delta_s + \gamma_s] + \sum_{s=0}^{t} \delta_s g_s + \sum_{s=0}^{t-1} \varepsilon_{s+1} h_s + z_{it}$$

and

$$(13) \quad y_t = \sum_{s=0}^{t} [\alpha_s + \beta_s\delta_s + \gamma_s] + \sum_{s=0}^{t} \delta_s g_s + \sum_{s=0}^{t-1} \varepsilon_{s+1} h_s + h_t,$$

where

$$(14) \quad g_t \equiv g[\bar{z}_{1t} + \eta_{1t}, \ldots, \bar{z}_{mt} + \eta_{mt}],$$

$$h_t \equiv h(z_{1t}, \ldots, z_{mt}; \bar{z}_{1t} + \eta_{1t}, \ldots, \bar{z}_{mt} + \eta_{mt}),$$

$$(15) \quad z_{it} \equiv \sum_{s=0}^{t} [(1-\delta_s)\zeta_{is} + \theta_{is}] \times \left[\prod_{j=1}^{t-s} (1 - \delta_{t-j+1})(1 - \varepsilon_{t-j+1})\right],$$

and $\bar{z}_{it} \equiv (1 - \varepsilon_t) z_{it-1} + \zeta_{it}$.

The proof is in the Appendix. Proposition 1 expresses the outcome of the learning process as the cumulative effect of trend terms α_t, $\beta_t\delta_t$, and γ_t and shock terms $\delta_t g_t$ and $\varepsilon_{t+1} h_t$, the effects of which persist indefinitely.[19] This persistence makes the process resemble a random walk in the aggregate, but with nonzero drift and declining variances determined by the parameters and the treatment variables. As explained in Section I, subsection C, α_0 and the α_t represent the average level of players' initial beliefs and any trends in their beliefs, the β_t correct their bidding rules for the expected cost of strategic uncertainty, and the γ_t reflect their average forward induction inferences on observing q_t. The $\delta_t g_t$ reflect the effect of dispersion in players' bids on their beliefs via the market-clearing price q_t. The $\varepsilon_{t+1} h_t$ reflect the robustness effect of the dispersion of the winners' efforts via y_t, characterized in Broseta (1993a Proposition 3.1) and Crawford (1995 Proposition 1), modified here by the optimistic subjects effect. Proposition 1 is valid no matter how the shocks are generated, but much of its usefulness stems from the fact that the shock terms are known functions of the z_{it}, $\bar{z}_{it}$, and η_{it}, which are *ex ante* i.i.d. across i for any t. This and the scaling properties of $g(\cdot)$ and $h(\cdot)$ are crucial to the comparative dynamics analysis below.

C. *Convergence*

In the next proposition, for technical reasons, we bound players' efforts. This is done by increasing x_{it} to its lower bound, denoted $\underline{x}$, or reducing it to its upper bound, denoted $\bar{x}$, when it would otherwise fall outside the interval $[\underline{x}, \bar{x}]$.

PROPOSITION 2: *Assume that the* x_{it} *are bounded, so that they remain in the interval* $[\underline{x}, \bar{x}]$. *If* $\sum_{s=0}^{\infty} \alpha_s$, $\sum_{s=0}^{\infty} \beta_s$, $\sum_{s=0}^{\infty} \gamma_s$, $\sum_{s=0}^{\infty} \sigma_{\theta s}^2$, $\sum_{s=0}^{\infty} \sigma_{\eta s}^2$, *and* $\sum_{s=0}^{\infty} \sigma_{\zeta s}^2$ *are finite, and* $0 < \delta_t, \varepsilon_t \leq 1$ *for all t, with* $(1 - \varepsilon_t)(1 - \delta_t)$ *bounded below* 1 *for sufficiently large t, then*

[19] The remaining terms, z_{it} in (12) and h_t in (13), are subsumed in the shock terms after the period in which they first appear; and the remaining endogenous variables, $\bar{x}_{it}$, p_{it}, and q_t, are easily computed from x_{it-1} and y_{t-1} using (1), (7), and (9).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

q_t, y_t, and the $\bar{x}_{it}$ and x_{it} converge, with probability one, to a common, finite limit.

The proof is in the Appendix. Proposition 2 gives conditions under which the learning dynamics converge to the bids and efforts associated with one of the symmetric pure-strategy subgame-perfect equilibria of the stage game. The model's implications for equilibrium selection are best summarized by the prior probability distribution of that equilibrium, which by Proposition 1 is typically nondegenerate due to the persistent effects of strategic uncertainty. Although Proposition 2 is quite helpful in understanding the model, its conditions are for convergence with probability one, and are only sufficient; convergence can easily occur when they are violated. The variance conditions are as one would expect from the strong law of large numbers, and probably cannot be weakened significantly if one insists on convergence with probability one. However, the estimated model reported in Section III, subsection E, converges strongly to the upper bound $\bar{x}$ in VHBB's environment, even though our parameter estimates violate Proposition 2's conditions on the β_s and γ_s. This suggests that our conditions on the α_s, β_s, and γ_s could be weakened, but the additional insight this would yield seems unlikely to justify the effort.

D. *Qualitative Comparative Dynamics*

Proposition 1 shows that unless the behavioral parameters vary sharply with changes in the treatment variables, the dynamics and limiting outcome respond to such changes in stable, predictable ways. We now conduct a qualitative comparative dynamics analysis of changes in m, n, and $f(\cdot)$, holding the behavioral parameters constant.

We define the jth order statistic of an empirical distribution $(\xi_1, \ldots, \xi_n)$ as the jth smallest of the ξ_i; thus, the minimum is the first order statistic and the median (for odd n) is the $(n + 1/2)$th. If $f(\cdot)$ is an order statistic, the index j identifies which one; and if $f(\cdot)$ is a convex combination of order statistics, j indexes the order statistics from which it is computed. An "increase in j" refers to any shift in the weights used to compute $f(\cdot)$ that increases j in the sense of first-order stochastic dominance, and a random variable is said to "stochastically increase" if its probability distribution shifts upward in this sense.

PROPOSITION 3: *Holding m and n and the behavioral parameters constant, increasing j stochastically increases y_t and x_{it} for all i and t.*

PROOF:

The proof is immediate from Proposition 1, noting that for any history of shocks, increasing j increases h_t for all t and leaves the g_t unaffected.

PROPOSITION 4: *Holding m and j (or the weights on alternative values of j) and the behavioral parameters constant, increasing n stochastically decreases y_t and x_{it} for all i and t.*

PROOF:

The proof is immediate from Proposition 1, noting that for any history of shocks, increasing n (with i running from 1 to the larger value of n) decreases both g_t and h_t for all t.

PROPOSITION 5: *Holding n and j (or the weights on alternative values of j) and the behavioral parameters constant, increasing m stochastically increases y_t and x_{it} for all i and t.*

PROOF:

The proof follows from Proposition 1. It is clear that for any history of shocks, increasing m increases g_t. It is not true that increasing m increases h_t for any history, because such an increase may cause turnover among those who win the right to play. However, it follows from (8) and the fact that the θ_{it} are i.i.d. with mean zero that any turnover replaces players with players whose effort distributions first order stochastically dominate them (whether or not the x_{it} are bounded), which stochastically increases h_t.

Propositions 3 and 4 extend Broseta (1993a Lemma 3.1) and Crawford (1995 Propositions 3–4), making precise, in the probabilistic sense appropriate to the model, the intuitions that coordination is less efficient when more

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

efficient outcomes are robust only to deviations by smaller sets of players, and less efficient in larger groups because it requires coherence among a larger number of independent decisions. Proposition 5 addresses a new issue, showing that increased competition for the right to play favors efficiency because it tends to yield higher market-clearing prices and intensifies the optimistic subjects effect.

E. *Quantitative Comparative Dynamics*

We now conduct a quantitative comparative dynamics analysis, showing how the mean coordination outcome is affected by changes in the treatment variables and the behavioral parameters. The analysis requires some additional notation. Let σ_{pt}^2 and σ_{zt}^2, respectively, denote the common *ex ante* variances of the p_{it} (or, equivalently, the $\bar{z}_{it} + \eta_{it}$) and the z_{it}. Because the ζ_{it} and θ_{it} are serially independent, (15) and (7) imply that

$$(16)\quad \sigma_{zt}^2 \equiv \sum_{s=0}^{t} [(1-\delta_s)^2\sigma_{\zeta s}^2 + \sigma_{\theta s}^2] \times \left[\prod_{j=1}^{t-s}(1-\delta_{t-j+1})(1-\varepsilon_{t-j+1})\right]^2$$

and

$$(17)\quad \sigma_{pt}^2 \equiv (1-\varepsilon_t)^2\sigma_{z-1}^2 + \sigma_{\zeta t}^2 + \sigma_{\eta t}^2.$$

Finally, let $\mu_t \equiv Eg((\bar{z}_{1t} + \eta_{1t})/\sigma_{pt}, \ldots, (\bar{z}_{mt} + \eta_{mt})/\sigma_{pt})$ and $\nu_t \equiv Eh(z_{1t}/\sigma_{zt}, \ldots, z_{nt}/\sigma_{zt}; (\bar{z}_{1t} + \eta_{1t})/\sigma_{pt}, \ldots, (\bar{z}_{mt} + \eta_{mt})/\sigma_{pt})$. Because the random variables $(\bar{z}_{it} + \eta_{it})/\sigma_{pt}$ and z_{it}/σ_{zt} are standardized, with mean 0 and variance 1, μ_t and ν_t are purely statistical parameters, completely determined by the joint distribution of these random variables and the treatment variables m, n, and $f(\cdot)$; they are subscripted because that distribution is time-dependent, and their dependence on m, n, and $f(\cdot)$ is suppressed for clarity. μ_t and ν_t respond to changes in m, n, and $f(\cdot)$ in the expected directions, as can be inferred from the responses of g_t and h_t identified in the proofs of Propositions 3–5.

PROPOSITION 6: *The ex ante means of y_t and the x_{it} are given, for all i and t, by*

$$(18)\quad Ex_{it} = \sum_{s=0}^{t}[\alpha_s + \beta_s\delta_s + \gamma_s] + \sum_{s=0}^{t}\delta_s\sigma_{ps}\mu_s + \sum_{s=0}^{t-1}\varepsilon_{s+1}\sigma_{zs}\nu_s$$

and

$$(19)\quad Ey_t = \sum_{s=0}^{t}[\alpha_s + \beta_s\delta_s + \gamma_s] + \sum_{s=0}^{t}\delta_s\sigma_{ps}\mu_s + \sum_{s=0}^{t-1}\varepsilon_{s+1}\sigma_{zs}\nu_s + \sigma_{zt}\nu_t.$$

PROOF:
Using (2), (4), (5), and (14) yields

$$(20)\quad Eg_s \equiv E[\sigma_{ps}g((\bar{z}_{1s} + \eta_{1s})/\sigma_{ps}, \ldots, (\bar{z}_{ms} + \eta_{ms})/\sigma_{ps})] \equiv \sigma_{ps}\mu_s$$

and

$$(21)\quad Eh_s \equiv E[\sigma_{zs}Eh(z_{1s}/\sigma_{zs}, \ldots, z_{ns}/\sigma_{zs}; (\bar{z}_{1s} + \eta_{1s})/\sigma_{ps}, \ldots, (\bar{z}_{ms} + \eta_{ms})/\sigma_{ps})] \equiv \sigma_{zs}\nu_s.$$

Taking expectations in (12) and (13) then yields (18) and (19).

Proposition 6 expresses the mean coordination outcome as a simple function of the behavioral parameters; the σ_{pt} and σ_{zt}, which can be computed from the behavioral parameters using (16) and (17); and the statistical parameters μ_t and ν_t, which reflect the influence of m, n, and $f(\cdot)$ on the dynamics. To apply this result we must estimate the behavioral parameters and evaluate μ_t and ν_t. The estimation is carried out in Section III. μ_t and ν_t are difficult to evaluate because their dependence on the distribution of the $\bar{z}_{it} + \eta_{it}$ and z_{it} is complex,

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

but it is possible to approximate them, as we now describe.

The approximations are based on further assumptions. The first addresses a difficulty in evaluating Eh_t that arises because unless $\sigma^2_{\eta t} = \sigma^2_{\theta t} = 0$, so that there is no dispersion in players' bidding rules and their inferences on observing q_t, the highest bidders need not be those who would choose the highest efforts in the coordination game. We assume that

$$(22)\quad Eh_t \equiv Eh(z_{1t}, \ldots, z_{mt}; \bar{z}_{1t} + \eta_{1t}, \ldots, \bar{z}_{mt} + \eta_{mt}) \approx Eh(z_{1t}, \ldots, z_{mt}; z_{1t}, \ldots, z_{mt}),$$

so that the increases and decreases in $h(\cdot)$ caused by such "crossovers" approximately cancel out on average. This should not be an important source of bias, because

$$(23)\quad z_{it} = (1 - \delta_t)\bar{z}_{it} + \theta_{it} = (1 - \delta_t)(\bar{z}_{it} + \eta_{it}) - (1 - \delta_t)\eta_{it} + \theta_{it},$$

so that there are no systematic differences between the orderings of the z_{it} and $\bar{z}_{it} + \eta_{it}$.

This approximation yields a simple characterization of the optimistic subjects effect, in which the order statistic is adjusted to reflect the auction's tendency to select players who choose higher efforts. In VHBB's environment, where the right to play a nine-person median game was auctioned in a group of 18, Eh_t is approximated by the robustness effect in an 18-person game without auctions in which the order statistic is the fifth highest effort (the median of the nine highest efforts). Thus, for the purpose of evaluating Eh_t the auctions effectively transformed VHBB's median game, which without auctions would have a robustness effect that contributes zero drift to the dynamics, into a 75th-percentile game ($0.75 = 13.5/18$) with a robustness effect that contributes a large upward drift.

To estimate the magnitude of this drift we assume, as in Broseta (1993a) and Crawford (1995), that $\bar{z}_{it} + \eta_{it}$ and z_{it} are approximately jointly normal for all i and t. Normality is a reasonable approximation because $\bar{z}_{it} + \eta_{it}$ and z_{it} are weighted sums of serially independent random variables. It makes the distributions of $(\bar{z}_{it} + \eta_{it})/\sigma_{pt}$ and z_{it}/σ_{zt} independent of t as well as i, so that $\mu_t \overset{t}{\equiv} \mu$ and $\nu_t \overset{t}{\equiv} \nu$. Given that $(\bar{z}_{it} + \eta_{it})/\sigma_{pt}$ and z_{it}/σ_{zt} are independent across i, μ and ν can be estimated [using (22) for ν] from D. Teichroew (1956 Table I) for any order statistic and $m \leq 20$. For VHBB's environment this yields $\mu = -0.069$ and $\nu = 0.665$, the respective means of the tenth and fifth largest [the $(n + 1)$st and the median of the n largest] of 18 i.i.d. standard normal variables.

Continuing to ignore the discreteness of effort, Proposition 6 can now be used to analyze the sources of the efficiency-enhancing effect of auctions in VHBB's experiment. Assume, as in Section III's econometric analysis, that $\varepsilon_t = \varepsilon$ for $t = 1, \ldots$, and $\delta_t = \delta$ for $t = 0, \ldots$, and set $\alpha_0 = \bar{\alpha}$ and $\alpha_t = 0$ for $t = 1, \ldots$.[20] Proposition 6 then implies

$$(24)\quad Ex_{it} = \bar{\alpha} + \sum_{s=0}^{t} [\delta\beta_s + \gamma_s] + \delta\mu \sum_{s=0}^{t} \sigma_{ps} + \varepsilon\nu \sum_{s=0}^{t-1} \sigma_{zs}$$

and

$$(25)\quad Ey_t = \bar{\alpha} + \sum_{s=0}^{t} [\delta\beta_s + \gamma_s] + \delta\mu \sum_{s=0}^{t} \sigma_{ps} + \varepsilon\nu \sum_{s=0}^{t-1} \sigma_{zs} + \nu\sigma_{zt},$$

which express Ex_{it} and Ey_t as functions of μ, ν, and identified behavioral parameters. $\bar{\alpha}$ is the average initial level of players' beliefs; $\delta\beta_s$

[20] Setting $\alpha_t = 0$ involves no loss of generality because, as (26)–(27) make clear, of the four parameters in the constant terms of (18) and (19), $\alpha_t + \beta_t\delta_t + \gamma_t$, only δ_t and the combinations $\alpha_t + \beta_t$ and $(1 - \delta_t)\alpha_t + \gamma_t$ are identified. Setting $\alpha_t = 0$ replaces those combinations by β_t and γ_t, which now partly reflect the trends in beliefs in α_t.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

TABLE 1—APPROXIMATIONS OF THE DYNAMICS FOR VHBB'S EXPERIMENT

t	$\bar{\alpha}$	$\delta\beta_t$	γ_t	$\delta\mu\sigma_{pt}$	$\varepsilon\nu\sigma_{zt-1}$	$\nu\sigma_{zt}$	Ex_{it}	Ey_t	$\overline{Ex_{it}}$	$\overline{Ey_t}$
0	4.940	−0.050	0.297	−0.036	—	0.811	5.151	5.962	5.542	5.750
1	—	−0.060	0.358	−0.025	0.622	0.554	6.046	6.600	6.166	6.375
2	—	−0.068	0.400	−0.020	0.427	0.444	6.785	7.000	6.610	6.625
3	—	−0.073	0.432	−0.017	0.341	0.379	7.000	7.000	6.721	6.875
4	—	−0.077	0.458	−0.015	0.291	0.335	7.000	7.000	6.902	6.875

reflects the adjustment of players' bids for strategic uncertainty and any trends in beliefs; and γ_s combines the forward induction effect and trends in beliefs. $\delta\mu\sigma_{ps}$ is the effect of the dispersion of bids via q_t. $\varepsilon\nu\sigma_{zs}$ is the persistent component of the optimistic subjects/robustness effect, which incorporates the inertia of players' adjustments through ε. Finally, $\nu\sigma_{zt}$ in (25) is the direct effect of the dispersion of efforts on Ey_t. Part of this last effect is transient; after the period in which it first appears it is multiplied by ε and incorporated into the sum in the preceding term.

Table 1 presents approximations of Ex_{it} and Ey_t based on (24) and (25) and Section III's estimates of the behavioral parameters, reported in Table 2. They are computed under the parameter restrictions described above and simple intertemporal restrictions on β_t, γ_t, and the shock distributions imposed in estimation, described in Section III, subsection E. The second through sixth columns in Table 1 give the terms in (24) and (25) for periods 0–4, with the implied values of Ex_{it} and Ey_t and the average values of x_{it} and y_t in the eight trials of VHBB's experiment, $\overline{Ex_{it}}$ and $\overline{Ey_t}$, for comparison. Ex_{it} is computed by summing the second through fifth columns in row t and adding the result to Ex_{it-1}, truncating at 7.000 as necessary. (For Ex_{i0}, Ex_{it-1} is replaced by the estimated value of $\bar{\alpha}$, 4.940.) Ey_t is computed by adding the transient term $\nu\sigma_{zt}$ to Ex_{it} and truncating at 7.000.

Because some parameters are not precisely estimated, we view Table 1's results as illustrative. The approximations track the dynamics reasonably well, somewhat overestimating the rate at which Ex_{it} and $Ey_t \rightarrow 7$. They suggest that of the four components of Ex_{it} and Ey_t in our theoretical analysis, only two were important in VHBB's experiment: γ_t, forward induction, and $\varepsilon\nu\sigma_{zt-1}$, the optimistic subjects/robustness effect. Both favor efficiency, with roughly equal strength. In the three periods (0–2) it takes Ex_{it} and Ey_t to reach 6.785 and 7.000 from the initial level set by $\bar{\alpha} = 4.940$, γ_t contributes 1.049 to their increases, and $\varepsilon\nu\sigma_{zt-1}$ contributes 1.055. The transient term $\nu\sigma_{z2}$ contributes an additional 0.444 to the increase in Ey_t.[21]

Our methods can also be used to assess the likely effects of changes in the treatment variables, on the assumption that the behavioral parameters do not vary sharply with such changes. Consider replacing VHBB's median game with a nine-person minimum game, with all else unchanged. $\mu = -0.069$ as before because m and n are unchanged, but ν is now approximately the robustness effect in an 18-player game without auctions, in which payoffs and best responses are determined by the ninth largest (the minimum of the nine largest) of players' efforts. For this game Teichroew's Table I yields $\nu = 0.069$, as the optimistic subjects effect is almost neutralized by the robustness effect of the minimum game. This tenfold fall in ν reduces $\varepsilon\nu\sigma_{zt-1}$ and $\nu\sigma_{zt}$ in Table 1 proportionately, yielding approximate

[21] These contributions sum to more than 7 − 4.940 because the other two components contribute −0.259, and the approximated Ey_2 must be rounded down to 7.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

TABLE 2—PARAMETER ESTIMATES

t	$\bar{\alpha}$	β_t	γ_t	λ_T	δ	ε	$\sigma^2_{\nu t}$	$\sigma^2_{\zeta t}$	λ_V	ρ
0	4.940	−0.0949	0.2973	—	0.5287	—	1.4872	1.0000	—	0.5226
	(0.511)	(0.154)	(0.221)		(0.137)		(0.383)	—		(0.217)
1	—	−0.1144	0.3583	−0.2689	0.5287	0.9348	0.6938	0.4665	1.0999	0.5226
		(0.177)	(0.181)	(0.402)	(0.137)	(0.309)	(0.132)	(0.072)	(0.223)	(0.217)
2	—	−0.1276	0.3995	−0.2689	0.5287	0.9348	0.4442	0.2987	1.0999	0.5226
		(0.195)	(0.156)	(0.402)	(0.137)	(0.309)	(0.090)	(0.073)	(0.223)	(0.217)
3	—	−0.1378	0.4317	−0.2689	0.5287	0.9348	0.3237	0.2177	1.0999	0.5226
		(0.211)	(0.144)	(0.402)	(0.137)	(0.309)	(0.075)	(0.067)	(0.223)	(0.217)
4	—	−0.1464	0.4583	−0.2689	0.5287	0.9348	0.2533	0.1703	1.0999	0.5226
		(0.224)	(0.144)	(0.402)	(0.137)	(0.309)	(0.067)	(0.061)	(0.223)	(0.217)
5	—	−0.1537	0.4814	−0.2689	0.5287	0.9348	0.2072	0.1393	1.0999	0.5226
		(0.236)	(0.152)	(0.402)	(0.137)	(0.309)	(0.061)	(0.056)	(0.223)	(0.217)

Note: Asymptotic standard errors in parentheses; log $L = -834.955$.

values for Ey_t of 5.235, 5.629, 6.031, 6.448, and 6.879 in periods 0–4.

These values are consistently higher than those observed in Weber's (1994) single trial in this environment, which yielded y_t of 3, 4, 1, 1, and 3 in periods 0–4. We suspect the discrepancy arises because the minimum is too far from the median to justify ignoring changes in the behavioral parameters, and because the approximations ignore the bounds on x_{it} and y_t. To predict outcomes accurately by our methods in this new environment would require empirical information about behavior in nearby environments.

F. *Analysis and Simulations with Discrete Efforts*

Most of our analysis extends to the model with discrete efforts $\hat{x}_{it}$ outlined in Section I, subsection D. The dynamics are still Markov, with state vector $\bar{\mathbf{x}}_t$ representing players' beliefs; the only difference is in how beliefs determine efforts and y_t. Adapting the arguments in Crawford (1995 Sec. 5) shows that Propositions 0 and 2–5 hold as stated, and Propositions 1 and 6 hold approximately when the grid of feasible efforts is fine enough.

The prior probability distributions of the dynamics and limiting outcome can be estimated precisely for the model with discrete efforts only by repeated simulation. As a check on the approximations we computed frequency distributions for 5,000 simulation runs with the estimated parameters reported in Table 2, under normality assumptions described in Section III, subsection A. The simulation results are reported in Tables 3 and 4 and Figure 1. Comparing the simulations with the approximations in Section II, subsection E, suggests that the approximations overstate the rate of increase of $\hat{x}_{it}$ and y_t because they ignore their upper bounds, which cut off the upper tails of the normal distributions on which our estimates of μ_t and ν_t are based. Although this problem will probably arise whenever either an upper or a lower bound comes disproportionately into play, the analysis of Crawford (1995) suggests that our approximation technique will otherwise yield reasonably accurate results.[22]

[22] Crawford and Broseta (1995 Figures II–IV) report simulation estimates of Eq_t and Ey_t for alternative environments in which $f(\cdot)$ ("j" in Section II, subsection D) varies from the fifth to the third or seventh order statistic, n varies from 9 to 7 or 11, and m varies from 18 to 16 or 20, holding other treatment variables constant at VHBB's values. These results suggest that, at least for moderate

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

TABLE 3—MEAN VALUES OF q_t AND y_t ($t = 0, \dots, 9$)

t	Actual mean q_t, y_t	Predicted mean q_t, y_t
0	4.775, 5.750	4.782, 5.648
1	5.650, 6.375	5.454, 6.207
2	6.213, 6.625	6.465, 6.864
3	6.450, 6.875	6.768, 6.963
4	6.700, 6.875	6.921, 6.993
5	6.775, 7.000	6.979, 6.998
6	6.850, 7.000	6.996, 6.999
7	6.838, 7.000	6.999, 7.000
8	6.913, 7.000	7.000, 7.000
9	6.988, 7.000	7.000, 7.000

III. Econometric Specification and Estimation

This section provides a general econometric framework for estimating the behavioral parameters and reports estimates and tests for VHBB's (1993) experimental data.

A. The Model

We continue to assume that players' bids, p_{it}, the market-clearing price, q_t, and their beliefs, x_{it}, are continuously variable, but we now assume that players' efforts are determined by an ordered probit model of discrete choice, described in Section III, subsection B.

Substituting $\bar{x}_{it}$ from (9) into (7)–(8) yields

$$(26) \quad p_{it} = \beta_t + \alpha_t + \varepsilon_t y_{t-1} + (1 - \varepsilon_t)x_{it-1} + \zeta_{it} + \eta_{it}$$

and

$$(27) \quad x_{it} = \gamma_t + (1 - \delta_t)\alpha_t + \delta_t q_t + (1 - \delta_t)\varepsilon_t y_{t-1} + (1 - \delta_t)(1 - \varepsilon_t)x_{it-1} + (1 - \delta_t)\zeta_{it} + \theta_{it}$$

for $t = 1, \dots,$ and $i = 1, \dots, m$, with initial conditions given by setting presample values to 0. For simplicity we write $\tau_{it} \equiv \zeta_{it} + \eta_{it}$, and $\nu_{it} \equiv (1 - \delta_t)\zeta_{it} + \theta_{it}$. We maintain the following distributional assumptions. Conditional on $\mathbf{I}_t$, the history of p_{is} and x_{is} through time t (up to and including y_t), the vector of innovations (τ_{it}, ν_{it}) is jointly i.i.d. across i and serially uncorrelated, and has a bivariate normal distribution with means 0, variances $\sigma^2_{\tau t} \equiv \sigma^2_{\zeta t} + \sigma^2_{\eta t}$ and $\sigma^2_{\nu t} \equiv (1 - \delta_t)^2\sigma^2_{\zeta t} + \sigma^2_{\theta t}$, and covariance $(1 - \delta_t)\sigma^2_{\zeta t}$.

The model defined by (26)–(27) determines all m players' current bids and beliefs, as functions of the observed values of y_{t-1} and q_t and players' past beliefs x_{it-1} (but not their past efforts). For any given t, player i plays the coordination game if and only if (barring ties, which have zero probability) $p_{it} \geq q_t$. This yields two regimes for x_{it}: one in which $p_{it} < q_t$ and x_{it} is completely censored; and one in which $p_{it} \geq q_t$ and x_{it} is not directly observed, but acts as a latent variable determining player i's observed effort. Thus, sample separation is determined endogenously by whether or not $p_{it} < q_t$.[23]

The resulting model raises three econometric issues: discrete choice, endogenous sample selection, and latent explanatory variables. We discuss these in Section III, subsections B–D, constructing the likelihood function as we go. We then discuss specification and estimation issues and report estimates in Section III, subsection E.

changes, all three of these treatment variables have significant effects, close to those suggested by our approximations.

[23] We assume that the p_{it} are unobserved, and to make the computations tractable we treat q_t as exogenous. These are reasonable approximations in VHBB's treatment.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

TABLE 4—DYNAMICS OF y_t (t = 0, ... , 6)

	y_0	y_1	y_2	y_3	y_4	y_5	y_6
7	0.125	0.625	0.625	0.875	0.875	1.000	1.000
	0.055	0.329	0.654	0.871	0.964	0.994	0.999
6	0.500	0.125	0.375	0.125	0.125	0.000	0.000
	0.550	0.551	0.315	0.122	0.036	0.006	0.001
5	0.375	0.250	0.000	0.000	0.000	0.000	0.000
	0.382	0.117	0.031	0.007	0.000	0.000	0.000
4	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	0.013	0.002	0.000	0.000	0.000	0.000	0.000
≤3	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: Actual frequency distributions in first lines, predicted distributions in second lines.

B. *Discrete Choice*

Because players' efforts are naturally ordered by their payoff implications, we use an ordered probit discrete-choice model, as in Daniel McFadden (1984), in which the x_{it} determine players' bids and efforts. We partition the state space of x_{it}—the extended real line—into response regions C_k ($k = 1, \ldots, K$), where C_k is the interval $(a_{k-1}, a_k]$, and $a_0 = -\infty < a_1 < \cdots < a_K = +\infty$. We then set $\hat{x}_{it} = k$ if and only if $x_{it} \in C_k$. Under our normality assumption, player i's choice probabilities at time t, conditional on $\mathbf{I}_{t-1}$, can be written

$$(28)\quad \Pr[x_{it} \in C_k | \mathbf{I}_{t-1}] = \int_{C_k} f_x(x_{it})\, dx_{it}$$

$$= \Phi[(a_k - E_{t-1}x_{it})/\sigma_{\nu_t}]$$

$$- \Phi[(a_{k-1} - E_{t-1}x_{it})/\sigma_{\nu_t}],$$

$$k = 1, \ldots, K,$$

where $f_x(\cdot)$ and $\Phi(\cdot)$ denote the marginal density of x_{it} and the standard normal cumulative distribution function, respectively.

It is commonly assumed that the thresholds a_k are independent of i and t. We make the stronger assumption that $a_k = k + {}^1\!/_2$ for $k = 1, \ldots, 6$, so that $\hat{x}_{it}$ is determined simply by rounding x_{it} to the nearest integer in the set $\{1, \ldots, 7\}$.[24]

C. *Endogenous Sampling*

Because the subsample of n players whose efforts are observed each period is endogenously determined, and τ_{it} and ν_{it} are generally correlated for any given player, standard maximum-likelihood estimates of the parameters in (26)–(27) will suffer from sample selection bias unless corrective action is taken.[25]

Let Γ_t denote the set $\{i | p_{it} \geq q_t\}$ of players who (barring ties) play the coordination game at time t, which VHBB recorded along with the other data from their experiment. We

[24] There is no loss of generality in identifying the scales on which $\hat{x}_{it}$ and x_{it} are measured. The equal spacing of thresholds is restrictive, but it is optimal in VHBB's median game when players are risk neutral and x_{it} is player i's estimate of $E[y_t|\mathbf{I}_{t-1}]$ (Crawford 1995 footnote 27 p. 127), and fixing the a_k greatly reduces the number of parameters to be estimated and identifies the conditional variance of ν_{it}.

[25] Various corrective methods have been proposed in the literature. See for example G. S. Maddala (1983 Ch. 9) for limited dependent variable models and Charles Manski and McFadden (1981) or Takeshi Amemiya (1985 Ch. 9) for qualitative response models.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

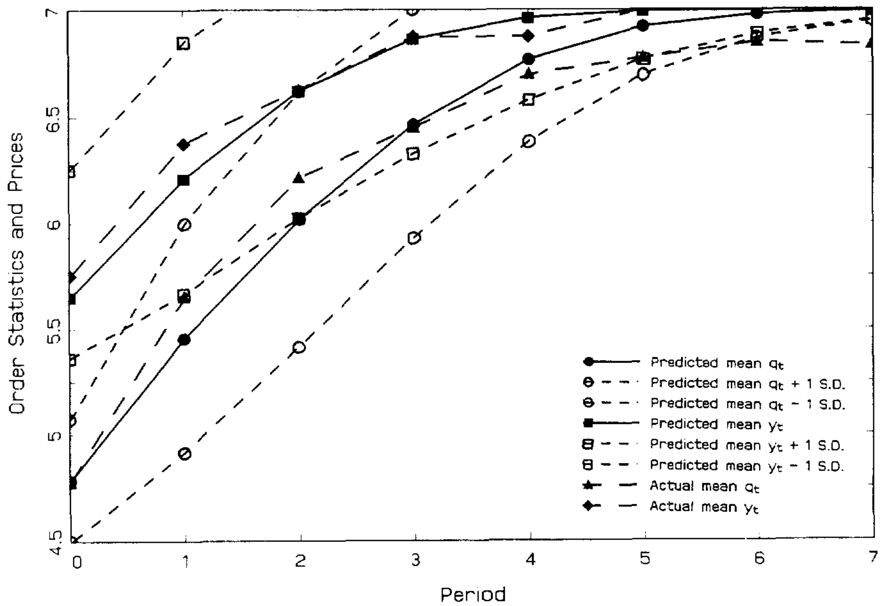

FIGURE 1. SIMULATED DYNAMICS

propose a full-sample estimation procedure, in which (26)–(27) are estimated simultaneously using all available information on subjects' decisions. For each $k = 1, \ldots, 7$, define the dummy variable $D_{it}^k \equiv 1$ when $\hat{x}_{it} = k$ and 0 otherwise. Let $\boldsymbol{\omega}_t$ denote the vector of parameters of interest at time t. Conditional on $\mathbf{I}_{t-1}$, the contribution of the (it)th observation to the likelihood function can then be written

$$(29)\quad L_{it}(\boldsymbol{\omega}_t) = 1_{\{i \notin \Gamma_t\}} \int_{-\infty}^{q_t} f_p(p_{it})\, dp_{it} + 1_{\{i \in \Gamma_t\}} \sum_{k=1}^{7} \left[\int_{C_k} \int_{q_t}^{\infty} f(x_{it}, p_{it})\, dx_{it}\, dp_{it} \right]^{D_{it}^k},$$

where $1_{\mathcal{A}}$ is the indicator function of the set $\mathcal{A}$, and $f(\cdot)$ and $f_p(\cdot)$ denote the joint conditional density of (x_{it}, p_{it}) and the marginal density of p_{it}, respectively. The two regimes in (29) reflect observations in which subject i does not play and we learn only that $p_{it} < q_t$, and those in which subject i plays and we learn both that $p_{it} \geq q_t$ and that $x_{it} \in C_k$ when $\hat{x}_{it} = k$. In this respect our model resembles differential response models for analyzing cross-section data from surveys with voluntary participation (Stephen Pudney, 1989 Sec. 2.6) or James Heckman's (1974) model of labor supply, with the added complications of the dynamic structure and the double integral in (29).

This estimation technique makes great demands on the sample. Both the sample selection equation, (26), and the equation describing the adjustment of beliefs, (27), must be estimated from observations of $\hat{x}_{it}$, y_t, and q_t; and q_t, for which we have only eight observations per period, plays a crucial role in identifying the parameters. Muthén and Joreskog's results, reported in Maddala (1983 p. 267), suggest that we may not obtain very precise estimates of the parameters $\beta_t + \alpha_t$, ε_t, and $\sigma_{\tau_t}^2$ in (26).

D. *Latent Explanatory Variables*

Standard estimation techniques, such as two-stage probit and Tobit models (David Grether and Maddala [1982] and Maddala [1983 Ch. 9]), deal with the latent variable

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

x_{it-1} on the right-hand sides of (26) and (27) by estimating the reduced form of the model. Here, however, the error terms of the reduced form are serially correlated, so that full maximum-likelihood estimation would involve calculating the T-fold integral of the joint distribution of x_{it} and p_{it} ($i = 1, \dots, m$; $t = 0, \dots, T$). This is computationally intractable even for simple types of serial correlation, and our model already involves additional difficulties due to endogenous sampling and a nonstationary error structure.

We address these problems by using the reduced forms of (26)–(27) to maximize the likelihood function $L(\boldsymbol{\omega}) \equiv \Pi_{i,t} L_{it}(\boldsymbol{\omega}_t)$, with $\boldsymbol{\omega} \equiv (\boldsymbol{\omega}_0, \dots, \boldsymbol{\omega}_T)$ and $L_{it}(\cdot)$ as in (29), with respect to the parameters of interest. Thus we ignore the serial correlation, which will result in inefficient parameter estimates, while taking the heteroskedasticity of the error terms fully into account as described below. We are not aware of any consistency or asymptotic normality results for maximum-likelihood estimators in models like ours. Thus our estimates of the behavioral parameters and standard errors must be interpreted with caution.

E. *Empirical Specification and Estimation*

As explained in Section I, subsection C, there is a danger of overfitting unless we rule out ad hoc time variation in the behavioral parameters. We therefore impose simplifying intertemporal restrictions on the parameters, which allow an adequate econometric description of behavior. As in Section II, subsection E, we set $\varepsilon_t = \varepsilon$ for $t = 1, \dots$ and $\delta_t = \delta$ for $t = 0, \dots$. We also set $\alpha_0 = \bar{\alpha}$ and $\alpha_t = 0$ for $t = 1, \dots$, so that $\alpha_t + \beta_t = \beta_t$ and $(1 - \delta)\alpha_t + \gamma_t = \gamma_t$ for $t = 1, \dots$; this involves no loss of generality because only $\alpha_t + \beta_t$ and $(1 - \delta)\alpha_t + \gamma_t$ are identified. We further assume that $\beta_t = \beta_0(1 + t)^{-\lambda_T}$ and $\gamma_t = \gamma_0(1 + t)^{-\lambda_T}$ for $t = 1, \dots$. Finally, following Broseta (1993b) and Crawford (1995), we impose the following intertemporal constraints on the variance-covariance matrix of the innovations τ_{it} and ν_{it} in (26)–(27): $\sigma^2_{\tau_t} = \sigma^2_{\tau_0}(1 + t)^{-\lambda_V}$, $\sigma^2_{\nu_t} = \sigma^2_{\nu_0}(1 + t)^{-\lambda_V}$, and $\mathrm{Cov}(\tau_{it}, \nu_{it}) = \mathrm{Cov}(\tau_{i0}, \nu_{i0})(1 + t)^{-\lambda_V}$, so that the correlation coefficient $\rho_t \equiv \rho$ for $t = 0, \dots$.

Under these assumptions, backward substitution for x_{it-1} in (26)–(27) yields

$$
\begin{aligned}
(30)\quad p_{it} = {} & \beta_0(1 + t)^{-\lambda_T} + \gamma_0(1 - \varepsilon) \\
& \times \sum_{k=1}^{t-1} [(1 - \delta)(1 - \varepsilon)]^{k-1}(t - k + 1)^{-\lambda_T} \\
& + (1 - \varepsilon)^t(1 - \delta)^{t-1}[\gamma_0 + (1 - \delta)\bar{\alpha}] \\
& + \varepsilon \sum_{k=1}^{t} [(1 - \delta)(1 - \varepsilon)]^{k-1} y_{t-k} \\
& + \delta(1 - \varepsilon) \\
& \times \sum_{k=1}^{t} [(1 - \delta)(1 - \varepsilon)]^{k-1} q_{t-k} + \tau_{it} \\
& + (1 - \varepsilon) \sum_{k=1}^{t} [(1 - \delta)(1 - \varepsilon)]^{k-1} \nu_{it-k}
\end{aligned}
$$

and

$$
\begin{aligned}
(31)\quad x_{it} = {} & \gamma_0 \sum_{k=0}^{t-1} [(1 - \delta)(1 - \varepsilon)]^k \\
& \times (t - k + 1)^{-\lambda_T} \\
& + [(1 - \varepsilon)(1 - \delta)]^t[\gamma_0 + (1 - \delta)\bar{\alpha}] \\
& + \varepsilon(1 - \delta) \\
& \times \sum_{k=1}^{t} [(1 - \delta)(1 - \varepsilon)]^{k-1} y_{t-k} \\
& + \delta \sum_{k=0}^{t} [(1 - \delta)(1 - \varepsilon)]^k q_{t-k} \\
& + \sum_{k=0}^{t} [(1 - \delta)(1 - \varepsilon)]^k \nu_{it-k}.
\end{aligned}
$$

Letting $e_{it} \equiv \tau_{it} + (1 - \varepsilon) \sum_{k=1}^{t} [(1 - \delta)(1 - \varepsilon)]^{k-1}\nu_{it-k}$ and $u_{it} \equiv \sum_{k=0}^{t} [(1 - \delta)(1 - \varepsilon)]^k \nu_{it-k}$ represent the error terms in (30) and (31), it is clear that the e_{it} and u_{it} are serially correlated for any given i, but

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

uncorrelated across i. The relevant components of the variance-covariance matrix are now given, for $t = 0, \ldots,$ by

$$(32)\quad \mathrm{Var}(u_{it}) = \sigma^2_{\nu_0} \sum_{k=0}^{t} [(1-\delta)(1-\varepsilon)]^{2k} \times (t-k+1)^{-\lambda_V},$$

$$(33)\quad \mathrm{Var}(e_{it}) = \sigma^2_{\tau_0}(1+t)^{-\lambda_V} + (1-\varepsilon)^2\sigma^2_{\nu_0} \times \sum_{k=1}^{t} [(1-\delta)(1-\varepsilon)]^{2(k-1)} \times (t-k+1)^{-\lambda_V},$$

and

$$(34)\quad \mathrm{Cov}(e_{it}, u_{it}) = \mathrm{Cov}(\tau_{i0}, \nu_{i0})(1+t)^{-\lambda_V} + (1-\varepsilon)\sigma^2_{\nu_0} \times \sum_{k=1}^{t} [(1-\delta)(1-\varepsilon)]^{2k-1} \times (t-k+1)^{-\lambda_V}.$$

For each period there are eight observations of q_t and y_t and, counting both regimes in (29), 144 of $\hat{x}_{it}$. The model defined by (28)–(34) was estimated by maximum likelihood, using the data from the first six periods of VHBB's eight trials.[26] There are ten parameters of interest: $\bar{\alpha}$, β_0, γ_0, λ_T, δ, ε, $\sigma^2_{\tau_0}$, $\sigma^2_{\nu_0}$, λ_V, and ρ.[27] All ten are theoretically identified. However, $\sigma^2_{\tau_0}$ is only weakly identified, because [by (5)] when the p_{it} are unobserved the only observable consequence of an increase in their variance is an increase in the variance of q_t. In practice the estimated $\sigma^2_{\tau_0}$ "blows up" because when $n/m = 0.5$, as in VHBB's design, this makes the fit in the sample selection equation (26) approximately independent of the estimated β_0 and ε, which can then be chosen to maximize the rest of the likelihood function. We deal with this problem by estimating the model under the normalization $\sigma^2_{\tau_0} \equiv 1$, which is arbitrary but unavoidably so when the p_{it} are unobserved. Fortunately, as just explained, the level of the $\sigma^2_{\tau_t}$ (as opposed to their time trend, λ_V, which is precisely estimated) has little effect on the model's implications for observable variables.

Table 2 reports the estimated parameters, with asymptotic standard errors in parentheses. The estimates confirm that the model gives an adequate econometric description of VHBB's subjects' behavior. With the exception of λ_T, the estimate of which is slightly and insignificantly less than zero, they satisfy the restrictions suggested by the theory.[28] The estimated adjustment coefficients δ and ε are plausible, though the latter, at 0.93, is much higher than the 0.58 estimate for the analogous parameter in the analogous treatment without auctions (Crawford, 1995 Table I) and somewhat higher than our prior. The estimated γ_t are positive and quite large, also somewhat higher than our prior; together with our estimate of δ they imply a large forward induction effect, raising efforts nearly half a unit above the value suggested by q_t and prior beliefs, on average. The estimated β_t are negative as expected, small (recalling that bids are scaled like efforts) and insignificantly different from 0. The estimated $\sigma^2_{\nu_t}$ are significantly positive, so that constraints ruling out strategic uncertainty are strongly rejected. The estimated cor-

[26] Although the trials lasted 10–15 periods, convergence was so rapid that there was little sample variation after period 6 and none after period 7. We excluded the data in every period for subject 1 in trial 10, who paid a high price to play the coordination game in period 3 and then played the lowest effort. Preliminary estimates revealed that this subject's behavior had a large impact on the estimated $\sigma^2_{\nu_0}$. The program code and the results of intermediate computations are available on request.

[27] Ten may seem like a large number, but Section II's analysis shows that to explain VHBB's result the model must describe the dynamics of a heterogeneous population of learning rules. Our estimates strongly suggest that no simpler model can do this, and if tighter estimates are desired further experimental data can be produced at will.

[28] λ_T is naturally positive because it represents the rate of *decline* of β_t and γ_t, which should approach zero over time as strategic uncertainty is eliminated by experience. Recall that α_t has been normalized to zero, so that β_t and γ_t include any exogenous trends in beliefs that show up in players' bids (β_t) and/or efforts (γ_t).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

relation coefficient ρ is significantly positive, as expected because ζ_{it} enters both τ_{it} and ν_{it}, with positive weights. Finally, the estimated λ_V is significantly positive, confirming that there were substantial learning effects, and greater than one (though insignificantly so), so that Proposition 2's variance conditions for convergence are satisfied.[29]

We estimated the model's implications for the prior probability distributions of the dynamics and limiting outcomes by repeated simulation, as described in Section II, subsection F, taking the discreteness and boundedness of efforts fully into account. Table 3 and Figure 1 report the actual and predicted mean values of q_t and y_t, and Table 4 reports the predicted distributions of y_t in the second lines of each cell, with the actual frequency distributions in the first lines for comparison. Figure 1 illustrates the dynamic interplay between q_t and y_t that underlies the efficiency-enhancing effect, in which q_t exerts upward pressure on the distributions of the x_{it} and y_t via forward induction, y_t then exerts upward pressure on the distributions of the $\bar{x}_{it+1}$, the p_{it+1}, and q_{t+1}, and so on.[30]

The actual and predicted means and distributions seem very close, with the single exception of the distribution of y_1. Their closeness is difficult to judge by eye, because the frequency distributions for eight trials must be coarser than the predicted distributions. However, χ^2 tests of goodness of fit, conducted for the distribution of y_t each period, never come close to rejecting the hypothesis that the actual frequencies were drawn from the predicted distributions.[31] The relevant χ^2 statistics, each with six degrees of freedom, are 0.854 with p-value 0.991 in period 0; 5.999 with p-value 0.423 in period 1; 1.826 with p-value 0.935 in period 4; and 0.350, 0.057, 0.048, and 0.008 in periods 2, 3, 5, and 6. These results suggest that the model closely reproduces the dynamics of subjects' interactions.

IV. Conclusion

In this paper we propose, estimate, and analyze a model of the learning dynamics in a recent experiment by Van Huyck et al. (1993), whose subjects repeatedly played a coordination game with seven Pareto-ranked equilibria, with the right to play auctioned each period in a larger group. The auctions had a striking efficiency-enhancing effect, in that subjects invariably bid the market-clearing price to a level recoverable only in the efficient equilibrium and then converged to that equilibrium, although subjects who played the same coordination game without auctions always converged to inefficient equilibria. The efficiency-enhancing effect of auctioning the right to play suggests a new and potentially important way in which competition may promote efficiency. And its effectiveness in focusing subjects' expectations on desirable equilibria may eventually help to unify our understanding of preplay communication, via "cheap talk" as well as costly signaling.

VHBB's contribution is a good example of how carefully designed and conducted experiments can change the way we think about important economic problems. Although it conveys a powerful impression by itself, its power to inform analysis can only be fully realized by a theory that identifies the mechanism behind the efficiency-enhancing effect and provides a firm basis for generalization to other laboratory or field environments. Working with a model that is flexible enough to nest alternative explanations, as in our earlier analyses of equilibrium selection in VHBB's (1990, 1991) experiments in Broseta (1993a, 1993b) and Crawford (1995), we find that explaining VHBB's (1993) result requires a model of stochastic, history-dependent learning dynamics, in which interactions between

[29] Because $\lambda_T < 0$, Proposition 2's conditions on β_t and γ_t are violated. As explained in Section II, subsection C, convergence occurs despite these violations because the x_{it} are bounded.

[30] The upward pressure on the x_{it} and y_t from q_t comes mainly from an estimated $\gamma_t > 0$. q_{t+1} is not generally higher than y_t because the estimated $\beta_t < 0$ and $\varepsilon < 1$.

[31] These tests are not independent across periods because of the history dependence; we offer them only as a way to gauge how closely the model reproduces the results.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

nonnegligible amounts of strategic uncertainty and the learning dynamics have a persistent effect on the efficiency of the limiting outcome. Our analysis quantifies this effect, relating it to the treatment variables (the numbers of players in the coordination game and the auction, and the order statistic that determines the robustness of desirable equilibria) in a way that makes it possible to assess its likely magnitude in other environments. We find that auctions can be expected to enhance efficiency to some extent in a wide range of environments that share the basic ingredients of the mechanism we identify—our optimistic subjects, forward induction, and robustness effects. However, in other environments this effect may be too small to assure fully efficient coordination.

Our methodology offers some more general insights about the analysis of learning and equilibrium selection. Because the extent of strategic uncertainty cannot usefully be explained by theory alone, its persistent effect on the limiting outcome gives the analysis an empirical component, which we describe via exogenous behavioral parameters in our specification of players' learning rules. Much of our analysis is independent of the precise values of these parameters, but they do affect the quantitative results. We close the model, when necessary, by estimating them using the experimental data. This dependence on empirical parameters and the need to analyze stochastic, history-dependent learning dynamics are likely to be encountered in realistic models of equilibrium selection in other environments. Although our analysis does not immediately generalize beyond VHBB's environments, we hope it shows that such analyses need not be intractable, and that the theoretical and empirical understanding needed to address economic questions involving equilibrium selection and coordination may be achievable in other settings.

APPENDIX

PROOF OF PROPOSITION 0:

Under the stated conditions $\eta_{it} \equiv \theta_{it} \equiv \zeta_{it} \equiv 0$ for $t = T, \ldots$. (11) reduces to

$$\text{(A1)} \quad \bar{x}_{it+1} = \delta_t g(\bar{x}_{1t}, \ldots, \bar{x}_{mt}) + \varepsilon_{t+1}(1-\delta_t) \times h(\bar{x}_{1t}, \ldots, \bar{x}_{mt}; \bar{x}_{1t}, \ldots, \bar{x}_{mt}) + (1-\varepsilon_{t+1})(1-\delta_t)\bar{x}_{it}.$$

Because δ_t, ε_{t+1}, $g(\bar{x}_{1t}, \ldots, \bar{x}_{mt})$, and $h(\bar{x}_{1t}, \ldots, \bar{x}_{mt}; \bar{x}_{1t}, \ldots, \bar{x}_{mt})$ are the same for all players and $0 \leq (1-\varepsilon_{t+1})(1-\delta_t) < 1$, (A1) leaves the order of the $\bar{x}_{it}$ (weakly) unchanged over time. It follows that there exist players, j and k, such that $g(\bar{x}_{1t}, \ldots, \bar{x}_{mt}) = \bar{x}_{jt}$ and $h(\bar{x}_{1t}, \ldots, \bar{x}_{mt}; \bar{x}_{1t}, \ldots, \bar{x}_{mt}) = \bar{x}_{kt}$ for all $t = T, \ldots$. (A1) then implies

$$\text{(A2)} \quad \bar{x}_{jt+1} = \delta_t \bar{x}_{jt} + \varepsilon_{t+1}(1-\delta_t)\bar{x}_{kt} + (1-\varepsilon_{t+1})(1-\delta_t)\bar{x}_{jt} = [1-\varepsilon_{t+1}(1-\delta_t)]\bar{x}_{jt} + \varepsilon_{t+1}(1-\delta_t)\bar{x}_{kt}$$

and

$$\text{(A3)} \quad \bar{x}_{kt+1} = \delta_t \bar{x}_{jt} + \varepsilon_{t+1}(1-\delta_t)\bar{x}_{kt} + (1-\varepsilon_{t+1})(1-\delta_t)\bar{x}_{kt} = \delta_t \bar{x}_{jt} + (1-\delta_t)\bar{x}_{kt}.$$

When $0 < \delta_t, \varepsilon_t \leq 1$, (A2) and (A3) imply that $\bar{x}_{jt}$ and $\bar{x}_{kt}$ converge to a common value weakly between $\bar{x}_{jT}$ and $\bar{x}_{kT}$, which is completely determined by $\bar{x}_{jT}$, $\bar{x}_{kT}$, the δ_t, and the ε_t. It is clear from (A1) and (8) that the other $\bar{x}_{it}$ and x_{it} converge to the same value, and from (6)–(9), (11), (A2), and (A3) that if $T = 0$, $\bar{x}_{it} \equiv x_{it} \equiv \alpha_0$ for all i and $t = 0, \ldots$.

PROOF OF PROPOSITION 1:

As in Crawford (1995 Proposition 1), the proof is immediate by induction once the solution has been found. Here we give the method of constructing the solution, which is informative. Substituting from (1), (3), and (6)–(9) yields

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$\text{(A4)}\quad y_t - y_{t-1}$$
$$\equiv h(x_{1t}, \dots, x_{mt}; p_{1t}, \dots, p_{mt})$$
$$- h(x_{1t-1}, \dots, x_{mt-1};$$
$$p_{1t-1}, \dots, p_{mt-1})$$
$$= \alpha_t + \beta_t\delta_t + \gamma_t$$
$$+ \delta_t g[(1-\varepsilon_t)x_{1t-1} + \zeta_{1t} + \eta_{1t}, \dots,$$
$$(1-\varepsilon_t)x_{mt-1} + \zeta_{mt} + \eta_{mt}]$$
$$+ h[(1-\delta_t)(1-\varepsilon_t)x_{1t-1}$$
$$+ (1-\delta_t)\zeta_{1t} + \theta_{1t}, \dots,$$
$$(1-\delta_t)(1-\varepsilon_t)x_{mt-1}$$
$$+ (1-\delta_t)\zeta_{mt} + \theta_{mt}; p_{1t}, \dots, p_{mt}]$$
$$- (1-\varepsilon_t)h[x_{1t-1}, \dots, x_{mt-1};$$
$$p_{1t-1}, \dots, p_{mt-1}].$$

It is clear from (12) that z_{it} is the idiosyncratic component of x_{it}, from (9) and (15) that $\bar{z}_{it}$ is the idiosyncratic component of $\bar{x}_{it}$, and from (7) that $\bar{z}_{it} + \eta_{it}$ is the idiosyncratic component of p_{it}. The second part of (15) is immediate from (9), and (8)–(9) imply that

$$\text{(A5)}\qquad z_{it} = (1-\delta_t)(1-\varepsilon_t)z_{it-1} + (1-\delta_t)\zeta_{it} + \theta_{it},$$

which yields the first part of (15) by successive substitution. Using (2), (4)–(5), (18), and (A5), (A4) can be rewritten

$$\text{(A6)}\quad y_t - y_{t-1} = \alpha_t + \beta_t\delta_t + \gamma_t$$
$$+ \delta_t g(\bar{z}_{1t} + \eta_{1t}, \dots, \bar{z}_{mt} + \eta_{mt})$$
$$+ h(z_1, \dots, z_{mt}; \bar{z}_{1t} + \eta_{1t}, \dots, \bar{z}_{mt} + \eta_{mt})$$
$$- (1-\varepsilon_t)h(z_{1t-1}, \dots, z_{mt-1};$$
$$\bar{z}_{1t-1} + \eta_{1t-1}, \dots, \bar{z}_{mt-1} + \eta_{mt-1}).$$

Summing (A6) and substituting from (14) yields (13). To derive (12) from (13), note that

$$\text{(A7)}\quad x_{it} - y_t \equiv x_{it} - h(x_{1t}, \dots, x_{mt}; p_{1t}, \dots, p_{mt})$$
$$\equiv z_{it} - h(z_{1t}, \dots, z_{mt}; \bar{z}_{1t} + \eta_{1t}, \dots, \bar{z}_{mt} + \eta_{mt})$$
$$\equiv z_{it} - h_t.$$

This completes the proof.

PROOF OF PROPOSITION 2:

The proof follows the martingale convergence arguments of Nevel'son and Has'minskii (1973 Theorem 2.7.3). Substituting (1), (3), and (7)–(9) into (26)–(27) shows that $\mathbf{x}_t$ can be taken as the state vector instead of $\bar{\mathbf{x}}_t$, which is convenient because it is the x_{it}, not the $\bar{x}_{it}$, that are directly affected by the bounds and by the discreteness of effort. Define the Lyapunov function $V_t \equiv \Sigma_{i,j}(x_{it} - x_{jt})^2$, where the summation is taken over all $i, j = 1, \dots, m$.[32] Clearly, $V_t \geq 0$ for all $\mathbf{x}_t$, with $V_t = 0$ if and only if $x_{it} = x_{jt}$ for all i and j. Substituting from (27) and simplifying yields

$$\text{(A8)}\quad V_t = \sum_{i,j} [(1-\varepsilon_t)(1-\delta_t)(x_{it-1} - x_{jt-1})$$
$$+ (1-\delta_t)(\zeta_{it} - \zeta_{jt}) + \theta_{it} - \theta_{jt}]^2$$
$$= \sum_{i,j} [(1-\varepsilon_t)^2(1-\delta_t)^2(x_{it-1} - x_{jt-1})^2$$
$$+ 2(1-\varepsilon_t)(1-\delta_t)(x_{it-1} - x_{jt-1})$$
$$\times \{(1-\delta_t)(\zeta_{it} - \zeta_{jt}) + \theta_{it} - \theta_{jt}\}$$
$$+ \{(1-\delta_t)(\zeta_{it} - \zeta_{jt}) + \theta_{it} - \theta_{jt}\}^2].$$

[32] Nevel'son and Has'minskii assume that $V_t \to \infty$ as $\|\mathbf{x}_t\| \to \infty$, but they use this condition only to ensure that solution paths are bounded; we assume boundedness directly.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Taking expectations in (A8) conditional on $\mathbf{x}_{t-1}$ then yields

$$\text{(A9)}\quad E_{t-1}(V_t|\mathbf{x}_{t-1})$$
$$= (1-\varepsilon_t)^2(1-\delta_t)^2 \sum_{i,j} (x_{it-1} - x_{jt-1})^2$$
$$+ 2(1-\varepsilon_t)(1-\delta_t)$$
$$\times \sum_{i,j} E\{(1-\delta_t)(\zeta_{it}-\zeta_{jt}) + \theta_{it} - \theta_{jt}\}$$
$$\times (x_{it-1} - x_{jt-1})$$
$$+ \sum_{i,j} E\{(1-\delta_t)(\zeta_{it}-\zeta_{jt}) + \theta_{it} - \theta_{jt}\}^2.$$

The first term on the right-hand side of (A9) is plainly bounded below V_{t-1} for all $\mathbf{x}_{t-1}$ outside any given neighborhood of the set for which $V_{t-1} = 0$. Without the bounds the second term equals 0, and the third term converges to 0 with probability 1 because the finiteness of $\Sigma_{s=0}^{\infty} \sigma_{\zeta s}^2$ and $\Sigma_{s=0}^{\infty} \sigma_{\theta s}^2$ implies that $(1-\delta_t)^2\sigma_{\zeta t}^2 + \sigma_{\theta t}^2 \to 0$, which implies that $E[(1-\delta_t)\zeta_{it} + \theta_{it}]^2 \to 0$ with probability 1. Thus, without the bounds $\{V_t\}$ eventually becomes a nonnegative supermartingale, so that $V_t \to 0$ with probability 1. Because the bounds can never increase $E[(1-\delta_t)\zeta_{it} + \theta_{it}]^2$ or [by (A8)] $E_{t-1}(V_t|\mathbf{x}_{t-1})$, this is also true for the bounded version of the $\{V_t\}$ process. In either case, it follows that for all i and j, $(x_{it} - x_{jt}) \to 0$ with probability 1. Using (9), the strong law of large numbers, and the continuity of $h(\cdot)$ then shows that $(y_t - x_{it})$ and $(\bar{x}_{it} - \bar{x}_{jt}) \to 0$ with probability 1. y_t, $\bar{x}_{it}$, and x_{it} must then converge to a common limit because $\bar{\mathbf{x}}_t$ and $\mathbf{x}_t$ cannot (with positive probability) return infinitely often to a given point at which $\bar{x}_{it} \neq \bar{x}_{jt}$ or $x_{it} \neq x_{jt}$ for some i and j; and $\bar{\mathbf{x}}_t$ or $\mathbf{x}_t$ cannot (with positive probability) oscillate infinitely often between distinct points at which $\bar{x}_{it} = \bar{x}_{jt}$ or $x_{it} = x_{jt}$ for all i and j, because $\Sigma_{s=0}^{t} \alpha_s$, $\Sigma_{s=0}^{t} \beta_s$, $\Sigma_{s=0}^{t} \gamma_s$, $\Sigma_{s=0}^{t} \theta_{is}$, and $\Sigma_{s=0}^{t} \zeta_{is}$ also converge to finite limits (the last two with probability 1). The convergence of p_{it} and q_t to the same limit follows from the continuity of $g(\cdot)$ and the facts that $\beta \to 0$ as $t \to \infty$ and $\Sigma_{s=0}^{t} \eta_{is}$ converges with probability one.

REFERENCES

Amemiya, Takeshi. *Advanced econometrics.* Cambridge, MA: Harvard University Press, 1985.

Ben-Porath, Elchanen and Dekel, Eddie. "Signaling Future Actions and the Potential for Sacrifice." *Journal of Economic Theory,* June 1992, *57*(1), pp. 36–51.

Broseta, Bruno. "Strategic Uncertainty and Learning in Coordination Games." Department of Economics Discussion Paper No. 93-34, University of California, San Diego, August 1993a.

______. "Estimation of a Game-Theoretic Model of Learning: An Autoregressive Conditional Heteroskedasticity Approach." Department of Economics Discussion Paper No. 93-35, University of California, San Diego, August 1993b.

Cachon, Gerard and Camerer, Colin. "Loss Avoidance and Forward Induction in Experimental Coordination Games." *Quarterly Journal of Economics,* February 1996, *111*(1), pp. 165–94.

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

Crawford, Vincent. "Adaptive Dynamics in Coordination Games." *Econometrica,* January 1995, *63*(1), pp. 103–43.

______. "Theory and Experiment in the Analysis of Strategic Interaction," in David Kreps and Ken Wallis, eds., *Advances in economics and econometrics: theory and applications,* Vol. I, Econometric Society Monographs No. 27. Cambridge: Cambridge University Press, 1997, pp. 206–42.

Crawford, Vincent and Broseta, Bruno. "What Price Coordination? Auctioning the Right to Play as a Form of Preplay Communication." Department of Economics Discussion Paper No. 95-41, University of California, San Diego, November 1995.

Grether, David and Maddala, G. S. "A Time Series Model with Qualitative Variables," in M. Deistler, E. Furst, and G. Schwodiauer, eds., *Games, economic dynamics, and time*

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

series analysis. Vienna: Physica Verlag, 1982, pp. 291–305.

Heckman, James. "Shadow Prices, Market Wages, and Labor Supply." *Econometrica*, July 1974, *42*(4), pp. 679–94.

Kim, Yong-Gwan. "Evolutionary Analyses of Tacit Communication in Van Huyck, Battalio, and Beil's Game Experiments." *Games and Economic Behavior*, October 1996, *16*(2), pp. 218–37.

Maddala, G. S. *Limited-dependent and qualitative variables in econometrics*, Econometric Society Monographs No. 3. Cambridge: Cambridge University Press, 1983.

Manski, Charles and McFadden, Daniel. "Alternative Estimators and Sample Designs for Discrete Choice Analysis," in Charles Manski and Daniel McFadden, eds., *Structural analysis of discrete data with econometric applications*. Cambridge, MA: MIT Press, 1981, pp. 2–50.

McCabe, Kevin; Rassenti, Stephen and Smith, Vernon. "Auction Institutional Design: Theory and Behavior of Simultaneous Multiple-Unit Generalizations of the Dutch and English Auctions." *American Economic Review*, December 1990, *80*(5), pp. 1276–83.

McFadden, Daniel. "Econometric Analysis of Qualitative Response Models," in Zvi Griliches and Michael Intriligator, eds., *Handbook of econometrics*, Vol. II. New York: Elsevier Science Publishing, 1984, pp. 1396–456.

Milgrom, Paul and Weber, Robert. "A Theory of Auctions and Competitive Bidding." *Econometrica*, September 1982, *50*(5), pp. 1089–122.

Nevel'son, M. B. and Has'minskii, R. Z. *Stochastic approximation and recursive estimation*, Vol. 47, Translations of Mathematical Monographs. Providence, RI: American Mathematical Society, 1973.

Pudney, Stephen. *Modelling individual choice: The econometrics of corners, kinks, and holes*. Oxford: Blackwell, 1989.

Roth, Alvin and Erev, Ido. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior*, January 1995, *8*(1), pp. 164–212.

Schelling, Thomas. *The strategy of conflict*, 1st Ed. Cambridge, MA: Harvard University Press, 1960.

Teichroew, D. "Tables of Expected Values of Order Statistics and Products of Order Statistics for Samples of Size Twenty and Less from the Normal Distribution." *Annals of Mathematical Statistics*, June 1956, *27*(2), pp. 410–26.

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

______. "Strategic Uncertainty, Equilibrium Selection Principles, and Coordination Failure in Average Opinion Games." *Quarterly Journal of Economics*, August 1991, *106*(3), pp. 885–910.

______. "Asset Markets as an Equilibrium Selection Mechanism: Coordination Failure, Game Form Auctions, and Tacit Communication." *Games and Economic Behavior*, July 1993, *5*(3), pp. 485–504.

Weber, Roberto. "The Role of Alternative Market Institutions on Equilibrium Selection in Coordination Games." B.A. honors thesis, Texas A&M University, 1994.

Woodford, Michael. "Learning to Believe in Sunspots." *Econometrica*, March 1990, *58*(2), pp. 277–307.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.