



Adaptive Dynamics in Coordination Games

Vincent P. Crawford

Econometrica, Vol. 63, No. 1. (Jan., 1995), pp. 103-143.

Stable URL:

<http://links.jstor.org/sici?sici=0012-9682%28199501%2963%3A1%3C103%3AADICG%3E2.0.CO%3B2-6>

Econometrica is currently published by The Econometric Society.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/econosoc.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

ADAPTIVE DYNAMICS IN COORDINATION GAMES

BY VINCENT P. CRAWFORD¹

This paper proposes a model of the process by which players learn to play repeated coordination games, with the goal of understanding the results of some recent experiments. In those experiments the dynamics of subjects' strategy choices and the resulting patterns of discrimination among equilibria varied systematically with the rule for determining payoffs and the size of the interacting groups, in ways that are not adequately explained by available methods of analysis. The model suggests a possible explanation by showing how the dispersion of subjects' beliefs interacts with the learning process to determine the probability distribution of its dynamics and limiting outcome.

KEYWORDS: Equilibrium selection, coordination, learning, strategic uncertainty.

1. INTRODUCTION

IN ECONOMICS, COORDINATION PROBLEMS are usually modeled as noncooperative games with multiple Nash equilibria in which any Pareto-efficient strategy combination is an equilibrium, but players' strategy choices are optimal only when they are based on sufficiently similar beliefs about how the game will be played. Although such games have no incentive problems as these are normally characterized, playing them often involves real difficulties. Similar difficulties lie at the heart of many questions usually analyzed under the assumption that players can coordinate on any desired equilibrium. These include how incentive schemes should be structured; which outcomes can be supported by implicit contract in a long-term relationship; whether, and how, bargainers share the surplus from making an agreement; and the role of expectations in macroeconomics.

Convincing answers to such questions must go beyond the observation that if rational players have commonly known, identical beliefs, then those beliefs must be consistent with some equilibrium in the game. However, the traditional approach to analyzing games with multiple equilibria relies on refining Nash's notion of equilibrium until (ideally) only one survives, and traditional refinements do not accomplish this for coordination games. This suggests that players are unlikely to base their decisions entirely on deductions from rationality, and highlights the importance of gathering information from other sources about how coordination problems are solved.

¹ I am grateful to Brian Arthur, Antonio Cabrales, John Conlisk, Robert Engle, Daniel Friedman, Clive Granger, Jerry Hausman, Yong-Gwan Kim, Mark Machina, Robert Porter, Garey Ramey, Michael Rothschild, Larry Samuelson, Joel Sobel, Maxwell Stinchcombe, Glenn Sueyoshi, John Van Huyck, Halbert White, Peyton Young, and anonymous referees for helpful suggestions; to Bruno Broseta and Pu Shen for valuable advice and outstanding research assistance; to Ray Battalio, Richard Beil, and John Van Huyck for access to their experimental data; to the Santa Fe Institute and the Department of Economics, University of Canterbury (New Zealand) for their hospitality; and to the National Science Foundation for research support.

Perhaps the most important source of such information now available is the rapidly growing experimental literature on coordination. A number of recent studies—such as Banks, Plott, and Porter (1988); Cooper, DeJong, Forsythe, and Ross (1990); Isaac, Schmidtz, and Walker (1989); Roth and Schoumaker (1983); and Van Huyck, Battalio, and Beil (1990, 1991) (henceforth “VHBB”)—report experiments in which subjects repeatedly played simple coordination games, uncertain in most cases only about each other’s strategy choices. The results suggest that players’ initial beliefs in such games are normally widely dispersed; that this dispersion, which I shall call *strategic uncertainty*, makes players’ experience with analogous games an important determinant of their decisions; and that interactions between strategic uncertainty and the process of learning from experience can exert a strong and lasting influence on coordination outcomes.

The effects of strategic uncertainty show up especially clearly in VHBB (1990, 1991). The large strategy spaces and the variety of modes of interaction in their designs yielded remarkably rich dynamics, with persistent patterns of discrimination among equilibria emerging over time. These patterns varied systematically with the environment, in ways that are not adequately explained by available methods of analysis but which can be better understood, as I shall argue, by taking strategic uncertainty fully into account.

This paper presents a simple model of the learning process that suggests a unified explanation of the dynamics and patterns of discrimination VHBB observed. The model is a repeated game in which players adjust their strategies in the underlying coordination game in response to their experience. It implies that players normally converge to some equilibrium in the underlying game, as usually happens in coordination experiments. The question remains, which equilibrium? The model answers this question by showing how strategic uncertainty interacts with the learning process to determine the prior probability distribution of its outcome. When players’ beliefs are identical from the start, the outcome is completely determined by their initial responses to the underlying coordination game, and does not vary with the environment except as their initial responses do. But for realistic levels of strategic uncertainty, declining at realistic rates as players learn to predict how the game will be played, the distribution of the dynamics and limiting outcome varies with the environment much as it did in the experiments.

The model’s predictions are influenced (but not completely determined) by the differences in the sizes of the basins of attraction of the equilibria in these environments discussed in Crawford (1991). Maynard Smith’s (1982) notion of evolutionary stability, Harsanyi and Selten’s (1988) risk-dominance axiom, and the techniques used by Kandori, Mailath, and Rob (1993) and Young (1993a) to study the limiting outcomes of related adjustment processes also respond to those size differences, discriminating among equilibria in ways that in some respects resemble the patterns VHBB observed. The extent to which these alternative approaches help to explain their results is discussed below.

The model departs from traditional noncooperative game theory in three main ways. First, players' behavior is adaptive rather than rational, in that they view their strategies in the underlying coordination game as the objects of choice, adjusting them over time in ways that are sensible but not necessarily consistent with equilibrium in the repeated game that describes the entire adjustment process or the underlying game.² One could conduct a traditional equilibrium analysis of the repeated game, but such an analysis would share the indeterminacy of the underlying game but nonetheless require the assumption that players' beliefs are coordinated when play begins. This begs the question of how coordination comes about, and often leads in practice to dictating a solution by applying refinements that are largely insensitive to the difficulty of coordination. It seems better for my purposes to allow players' beliefs to differ and study the dynamics of the process by which they converge.

Second, the model relies on simplifying assumptions about the structure of the strategic environment in the spirit of evolutionary game theory. These assumptions are satisfied by VHBB's designs (see Crawford (1991)) and are common to a number of interesting economic models (see Woodford (1990) and Cooper and John (1988)). They allow a more informative analysis of the implications of strategic uncertainty than now seems possible for more general games. The results are applicable to some important questions, and the methods seem likely to be useful in other settings.

Finally, instead of fully endogenizing players' beliefs and strategy choices (whose differences cannot be traced to differences in players' information or other characteristics) the model characterizes them statistically, treating certain aspects of the process that describes how they evolve as exogenous parameters to be estimated on a case-by-case basis. This departure reflects the conviction, first expressed by Schelling (1960), that it is impossible to predict from rationality alone how people will respond to coordination problems. I illustrate the usefulness of this approach by using the data from VHBB's experiments to estimate the model, treating the idiosyncratic components of players' beliefs as error terms, and then using the estimates to infer the prior probability distributions of outcomes in the various environments. In each case the model provides an adequate statistical summary of subjects' behavior while closely reproducing the dynamics of their interactions. The distributions it implies suggest that the observed dynamics are not anomalous, but in some cases they strengthen or modify the impressions created by the raw data.

The paper is organized as follows. Section 2 introduces the model and summarizes VHBB's experimental designs and results. Sections 3 and 4 carry out an analysis under the simplifying assumption that players' strategies are continuously variable. Section 5 extends the analysis to the case in which players

² Because players' strategy choices normally converge to an equilibrium in the underlying game, this "irrationality" does not persist, and is therefore immune to the most common criticism of adaptive analyses.

have discrete, finite strategy spaces, as in the experiments. Section 6 reports econometric estimates of the model's parameters based on VHBB's experimental data, and Section 7 discusses the implications of these estimates for the probability distributions of coordination outcomes.

2. VAN HUYCK, BATTALIO, AND BEIL'S EXPERIMENTAL DESIGNS AND RESULTS

In VHBB's (1990, 1991) experiments, subjects repeatedly played coordination games in which their payoffs were determined by their own strategies and simple summary statistics of other players' strategies. Explicit communication was not allowed, but the relevant summary statistic was publicly announced after each play. (In some experiments entire strategy profiles were also announced; this made little difference to the results.) With apparently unimportant exceptions, all details of the structure of each experimental environment were publicly announced at the start. There was ample evidence that the subjects understood the rules and were paid well enough to induce the desired preferences.³

VHBB's designs are best understood in simplified versions of the games used in their 1990 experiments. Suppose that a number of players choose between two efforts, 1 and 2. The minimum of their efforts determines their total output, which they share equally. Effort is costly, but sufficiently productive that if all players choose the same effort their output shares more than repay their costs. If anyone shirks, however, the balance of the others' efforts is wasted. Assuming for definiteness that output per capita is twice the minimum effort and each player's unit effort cost is 1, payoffs are determined as follows:

		Minimum Effort	
		2	1
Player's Effort	2	2	0
	1		1

A game with this structure was suggested by Bryant (1983) (see also Cooper and John (1988)) as a model of Keynesian effective demand failures. This game has a long history in economics, which can be traced to the stag hunt example Rousseau (1973 (1775), p. 78) used to discuss the origins of the social contract. To see the connection, imagine (adding some game-theoretic detail to Rousseau's discussion) that each of a number of hunters must independently decide whether to join in a stag hunt (effort 2) or hunt rabbits by himself (effort 1). Hunting a stag yields each hunter a payoff of 2 when successful, but success requires the cooperation of every hunter and failure yields 0. Hunting rabbits yields each hunter a payoff of 1 with or without cooperation, and thereby determines the opportunity cost of effort devoted to the stag hunt.

³ VHBB (1990, 1991) and Crawford (1991) describe the designs in more detail.

For any number of players, stag hunt has two pure-strategy Nash equilibria, one with all choosing effort 2 and one with all choosing effort 1. The effort-2 equilibrium is the best feasible outcome for all players, and all strictly prefer it to the effort-1 equilibrium. This rationale for playing effort 2 does not depend on game-theoretic subtleties; it is clearly the “correct” coordinating principle. However, effort 2’s higher payoff when all players choose it must be traded off against its risk of a lower payoff when someone does not. For a player to prefer effort 2 (treating the influence of his choice on future developments as negligible) he must believe that the correctness of this choice is sufficiently clear that it is more likely than not that all of the other players will believe that its correctness is sufficiently clear to all. Informal experiments suggest that people are often uncertain about whether other people will believe this, and that most people accordingly believe that effort 2 is a good bet in small groups but not in large groups.⁴

VHBB’s experiments showed that this kind of uncertainty can have profound and lasting consequences. Their subjects chose among seven efforts instead of two, with both the summary statistic used to determine payoffs and the size of the groups playing the game varying across treatments. In the “minimum” experiments reported in VHBB (1990) populations of 14–16 subjects repeatedly played games like stag hunt, first in large groups with the minimum effort in the entire population determining payoffs, and then in random pairs with each subject’s payoff determined by his current pair’s minimum.⁵ In the “median” experiments reported in VHBB (1991) populations of nine subjects repeatedly played games in which the entire population’s median effort determined their payoffs, with variations in the payoff function across three treatments. In each case the game had the “same” seven strict, symmetric, pure-strategy equilibria; this similarity in strategic structure extends to the experimental environments when the random-pairing treatment is viewed as a game played simultaneously by all members of the population, with players’ expected payoffs evaluated before the uncertainty of pairing is resolved.

These environments elicited roughly similar initial distributions of effort, with high to moderate means and variances. Subjects’ subsequent behavior, however, differed strikingly and systematically across treatments, with consequences for equilibrium selection that appeared likely to persist indefinitely. In the large-group minimum treatment subjects’ choices gravitated strongly toward the

⁴ Note that these beliefs are self-confirming. They are plausible because if players choose independently, with probabilities independent of the number of players, the clarity of the principle is less likely to be sufficient the larger the group. Because this intuition concerns a choice between strict equilibria, however, it is not captured by traditional refinements like trembling-hand perfectness or strategic stability.

⁵ In the random-pairing treatment each subject was told only his current pair’s minimum. There was also a treatment in which the minimum game was repeatedly played by fixed pairs, with subjects informed that their pairings were fixed. Those subjects were evidently aware of the importance of repeated-game strategies, and usually achieved efficient coordination. Because repeated-game strategies raise difficult new issues—the horizon over which players evaluate their strategies, for one—that treatment is not discussed here. It may be possible to explain its results along the lines suggested by Crawford and Haller (1990) or Kim (1990).

lowest effort, even though this led to the least efficient equilibrium. By contrast, in the random-pairing minimum treatment subjects' efforts converged very slowly with little or no trend; and in the median treatments subjects invariably converged to the equilibrium determined by the initial median, even though it varied across runs in each treatment and was usually inefficient.⁶ Thus, the dynamics were highly sensitive to group size, with very different rates of convergence and patterns of discrimination among equilibria in the random-pairing and large-group minimum treatments. They were also highly sensitive to the summary statistic used to determine payoffs, with no trends and very strong history-dependence in the median treatments but strong trends and no history-dependence in the large-group minimum treatment.

The similarity of the distributions of subjects' initial responses across treatments suggests that the differences in their subsequent behavior can be attributed to the dynamics of the learning process. Explaining VHBB's results requires imposing structure on this process. Many structures are consistent with some aspects of the results, and the dynamics will be at least as important as their limiting outcomes in discriminating among explanations. It is instructive, however, to consider the extent to which the outcomes that emerged (or seemed to be emerging) in the experiments can be explained by applying equilibrium refinements to the underlying coordination games.⁷

Traditional refinements like trembling-hand perfectness and strategic stability do not address the strategic issues raised by VHBB's games. Those refinements are motivated by asking whether a given equilibrium is self-enforcing if all players expect it to govern play. They therefore do not discriminate among the strict equilibria in these games, which all pose essentially the same strategic question in the absence of strategic uncertainty.

More promising are refinements like Harsanyi and Selten's (1988) risk-dominance axiom or the "general theory of equilibrium selection" of which it is a part, which do discriminate between strict equilibria. Although Harsanyi and Selten assume that players' beliefs always converge to a particular equilibrium before play begins, the mental tâtonnements that model their convergence process are sensitive to strategic uncertainty, with implications for equilibrium selection that are similar in some respects to the patterns VHBB observed. However, neither Harsanyi and Selten's theory nor a variant that eliminates the precedence they give payoff-dominance (allowing risk-dominance, which embodies most of their ideas about the effects of strategic uncertainty, to determine the outcome in most of VHBB's treatments) corresponds very closely to VHBB's

⁶ Isaac, Schmitz, and Walker (1989) and several more recent studies report similar results. Comparing the results across VHBB's (1990) treatments that differed only in subjects' prior experience suggests that the patterns they observed are not eliminated when subjects learn to anticipate them; if anything, they are reinforced.

⁷ Applying refinements to the repeated game that describes the entire learning process does not increase their explanatory power, and seems less plausible in this setting.

subjects' strategy choices.⁸ As will be seen, this is due mainly to the influence strategic uncertainty exerts on the learning dynamics, an influence their theory rules out by assumption.

Also more promising is Maynard Smith's (1982) notion of evolutionary stability. Formally a static equilibrium refinement, it is motivated by considering conditions for local stability of evolutionary dynamics that resemble simple models of the learning dynamics (see Crawford (1989)). Crawford (1991) explored the possibility of an "evolutionary" explanation of VHBB's results, verifying that VHBB's environments satisfy the structural assumptions of evolutionary game theory, so that the notion of evolutionary stability, suitably generalized to allow finite populations and interactions other than random pairing, can be applied. Like risk-dominance, evolutionary stability responds to the large differences in the sizes of basins of attraction in the large-group minimum treatment by discriminating among strict equilibria in a way that resembles VHBB's results. Although suggestive, an analysis along these lines does not determine players' initial strategy choices, as would be required to determine the patterns of equilibrium selection in VHBB's other treatments, and the analogy between evolution and learning may well be unreliable in environments with strategic uncertainty. The structure of evolutionary games nevertheless has important advantages in modeling the learning dynamics. Section 3 introduces a model that combines this structure with a flexible model of individual strategy adjustment, in which players learn from experience in comparatively sophisticated ways.

3. A MODEL WITH CONTINUOUSLY VARIABLE STRATEGIES

In this section and the next it is assumed that players' pure strategies are continuously variable. The analysis resolves most of the issues that arise when the model is extended to the discrete strategy spaces of VHBB's experiments in Sections 5–7.

⁸ As explained in Crawford (1991), Harsanyi and Selten's theory selects the equilibrium with effort 7 (in response to payoff-dominance) in all of VHBB's treatments but median treatment Φ , where it selects the equilibrium with effort 4 (in response to symmetry). These selections correspond to 7% of subjects' last-period choices in the large-group minimum treatment, 43% in the random-pairing minimum treatment, and 0%, 67%, and 33% in median treatments Γ , Ω , and Φ . These success rates are mostly better than random, which with seven effort levels would imply a 14% success rate, but low. Harsanyi and Selten's theory without payoff-dominance selects the equilibrium with effort 1 in the large-group minimum treatment (in response to risk-dominance); the equilibrium with effort 4 in the random-pairing minimum treatment (because risk-dominance is neutral in VHBB's two-person minimum game and the theory therefore applies the tracing procedure to a uniform prior over the undominated strategies); the equilibrium with effort 7 in median treatments Γ and Ω (in response to risk-dominance); and the equilibrium with effort 4 in median treatment Φ (in response to risk-dominance, after imposing symmetry). This raises the success rate to 72% in the large-group minimum treatment, lowers it to 17% in the random-pairing minimum treatment, and leaves it unchanged in the median treatments. These success rates are again better than random, but low. The patterns differ for first-period choices, but the success rates are no higher on average.

The model is defined for a class of environments with the structure of evolutionary games, which generalizes VHBB's experimental designs. There is a finite number of indistinguishable players, who repeatedly play a game in which their roles are symmetric. This game should be thought of as describing the interactions of the entire population of players, with their payoffs evaluated taking into account any uncertainty about how they interact. Players' strategies are identified, so that it is meaningful to say that different players choose the same strategy, or that a player chooses the same strategy in different periods. Players play only pure strategies, and their strategy spaces are one-dimensional.⁹

I focus on coordination games in which any symmetric combination of pure strategies is an equilibrium, and they are the only pure-strategy equilibria. Because players' roles are symmetric, these equilibria are Pareto-ranked unless players are indifferent between them. This makes the coordination problem particularly simple, with a one-dimensional continuum of possible limiting outcomes and an unambiguous measure of efficiency. I further restrict attention to games in which each player's best replies are given by a summary statistic of the strategies chosen by all players in the population whenever the statistic in question is unaffected by his choice. The summary statistic can be written in general as $y_t \equiv f(x_t)$, where $x_t = (x_{1t}, \dots, x_{nt})$ and x_{it} denotes player i 's strategy choice at time t . I assume that $f(\cdot)$ is continuous and, for any x_t and any 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, as in VHBB's minimum and median treatments, or a convex combination of order statistics such as the arithmetic mean.¹⁰ To see what they entail, note that because players' roles are symmetric $f(\cdot)$ must be a symmetric function of the x_{it} ; its value is therefore completely determined by the order statistics of their empirical distribution. My

⁹ The results of Crawford (1989) suggest that equilibria in which individuals play mixed strategies would not emerge if they were allowed. The polymorphic configurations of pure strategies that mixed-strategy equilibria often represent in evolutionary game theory are allowed, but do not emerge in the games studied here.

¹⁰ When players are risk-neutral their best replies in VHBB's random-pairing minimum treatment are actually given by the median effort in the population. A risk-neutral player i would like, ideally, to choose the x_i that maximizes

$$\begin{aligned} E[p \min\{x_i, x_j\} - qx_i] &= p[1 - G(x_i)]x_i + p \int_{-\infty}^{x_i} x_j dG(x_j) - qx_i \\ &= (p - q)x_i - p \int_{-\infty}^{x_i} G(x_j) dx_j, \end{aligned}$$

where $p > q > 0$, $G(\cdot)$ is the empirical distribution function of the other players' efforts, and the second equality is obtained by integrating the Stieltjes integral by parts. Computing the left- and right-hand derivatives of this objective function makes it clear that x_i^* solves this problem if and only if $G(x) \leq (\geq) 1 - q/p$ when $x < (\geq) x_i^*$, so that x_i^* is an order statistic of the distribution $G(\cdot)$. In the random-pairing minimum experiments $p = 2q$, so that x_i^* is the median when it is well-defined. VHBB also conducted some large-group minimum experiments in which the cost of effort was lowered to zero, making all efforts at or above the current minimum best replies and the highest effort a weakly dominant strategy. The effect of dominance on players' strategy choices is easily handled within the model, but the generic multiplicity of players' best replies raises new issues, not addressed here.

assumptions rule out most nonlinear functions of these order statistics, which is certainly restrictive but probably not unrepresentative of symmetric games.

The structure of the environment is made public before play begins, and the value of the relevant summary statistic is publicly announced after each play.¹¹ Players' beliefs and strategy choices evolve as follows. Their beliefs at the start of play determine their initial choices. They then observe the value of the summary statistic that results, update their beliefs and make new choices, and the process continues. Thus, players face uncertainty only about each other's strategy choices, and both the effects of those choices and the information they receive about them are filtered through the summary statistic. It follows that in choosing their strategies, they need form beliefs only about the summary statistic.

In this learning process, even if players form and revise their beliefs sensibly and choose their strategies optimally given their beliefs, their beliefs and strategy choices may differ in unpredictable ways. I imagine that each player begins with a prior about the process that generates y_t , and that each period he updates his prior in response to his new information and then chooses the strategy that is optimal given his beliefs about y_t . Players whose priors differ may then have different beliefs, even if they observe the same history and use the same procedures to interpret it. I suppose that each player treats the probability that his decisions influence the $\{y_t\}$ process as negligible, so that his optimal choice each period is determined by his beliefs about the current value of y_t ; this is not strictly necessary, but it is a natural simplification given the unimportance of such influences in VHBB's experiments.¹² Finally, I assume that each player's prior is sufficiently undogmatic that he will eventually learn to predict y_t correctly if it converges.

The importance of the differences in players' beliefs and the need to supplement the theory with empirical information about them make it essential to represent beliefs accurately and flexibly, and to describe their evolution in terms of observable variables so that it is possible to estimate the parameters left undetermined by the theory. It also seems advisable to avoid unduly restricting the form of players' priors about the $\{y_t\}$ process, given the lack of theory on this issue. These considerations suggest a model in the style of the adaptive control literature.¹³ The key insight of the control literature is that describing the evolution of agents' beliefs does not require that they be explicitly represented as probability distributions or their moments: It is enough to model the

¹¹ Although subjects were told only their pair minimum in the random-pairing minimum treatment, this can be viewed as a noisy estimate of the population median that determined their best replies (see footnote 10), so that the limitation on their information can be treated as an increase in the dispersion of their beliefs.

¹² In a large-group minimum game a player can expect to influence the minimum only if he reduces his effort below his estimate of the minimum of the others' choices, but the model implies that he will never do this. In more than 90% of the periods in the other treatments, no single subject's choice could have altered the summary statistic.

¹³ See for example Ljung and Soderstrom (1983) and Nevel'son and Has'minskii (1973). My discussion follows Woodford (1990, Section 2), who gives an excellent overview.

dynamics of the estimates of the optimal decisions they imply, which are the only aspects of agents' beliefs that affect the outcome. This simplification will accommodate the wide range of learning behavior VHBB's subjects exhibited and make it possible to describe the dynamics in a simple framework that satisfies the desiderata mentioned above.

The logic of this approach and the need for accuracy and flexibility suggest that beliefs are best represented by continuous variables. When strategies are also continuously variable it is possible, if players are risk-averse or risk-neutral, to represent their beliefs directly by their optimal decisions.¹⁴ Thus, in this section and the next, x_{it} will represent both beliefs and strategy choices. (In later sections x_{it} will continue to represent beliefs and remain continuously variable, with strategy choices determined by a standard model of discrete choice in which the x_{it} are latent variables.) Note that this way of representing players' beliefs requires no specific restrictions on their priors or risk preferences. There is also no need for a player's beliefs, as represented by his strategy choices, to be related in any simple way to the moments of the distribution that describes his beliefs about y_t . If, however, players come to expect a particular value of y_t the x_{it} will approach that value; thus, in the limit, their choices can be viewed as estimates of the mean of y_t .

I follow the control literature in assuming that players' beliefs and strategy choices evolve according to linear adjustment rules of the following form:

$$(1) \quad x_{i0} = a_{i0}$$

and

$$(2) \quad x_{it} = a_{it} + b_{it}y_{t-1} + (1 - b_{it})x_{it-1} \quad (t = 1, \dots).$$

The a_{it} and b_{it} in (1) and (2) are exogenous coefficients, which represent any trends in players' beliefs and how their beliefs respond to new information, thereby reflecting their precision. These coefficients are allowed to vary with i and t , as described below, to accommodate differences in players' beliefs and learning behavior.

Although (2) resembles partial adjustment to the strategy choice the latest observation of y_t suggests is optimal, it is taken here to represent full adjustment to his current estimate of his optimal decision, which normally responds less than fully to each new observation because it is only part of the information he has about the $\{y_t\}$ process.¹⁵ Suppose, for instance, that player i 's decisions are certainty-equivalent, so that his optimal strategy choice equals his current estimate of the mean of y_t . Then if he is convinced that the y_t are independent draws from a fixed distribution, and if he puts as much weight on his prior as he

¹⁴ When players' optimal decisions do not preserve enough information about their beliefs to represent them accurately, it is necessary to keep track of beliefs and decisions separately, as in the analysis of discrete strategy choice in Section 5.

¹⁵ Note that these learning rules differ from those discussed in the psychological learning literature (see Roth and Erev (1995) for an interesting application in economics) in that they incorporate information about the game's best-reply structure. This seems appropriate for VHBB's experiments, but might be unrealistic in other settings.

would on τ such draws, he will set $a_{it} \equiv 0$ and $b_{it} \equiv 1/(\tau + t + 1)$, as in "fictitious play." If he believes instead that the $\{y_t\}$ process is a driftless random walk, he will set $a_{it} \equiv 0$ and $b_{it} \equiv 1$, as in a "best-reply" process. Thus (2) includes as special cases two of the learning rules most often studied in the game theory literature.

In general players' priors are more complex than this, and their decisions may not be well approximated by certainty equivalence. With or without certainty equivalence or strong restrictions on priors, players whose adjustments follow (2) will learn to predict y_t if it converges and choose x_{it} that are optimal, given their predictions, as long as $a_{it} \rightarrow 0$ as $t \rightarrow \infty$ and $0 < b_{it} \leq 1$ for sufficiently large t . More generally, learning rules of this kind have been shown in the control literature to provide a robust, effective approach to the estimation problems faced by agents who understand the forecasting problems they face but are unwilling to make the specific assumptions about the process or unable to store and process the large amounts of information an explicitly Bayesian approach would require.

It is plain that the learning rules in (2) are usually less than fully rational in the sense used in game theory, because players' priors about the structure of the $\{y_t\}$ process are not required to be correct, and because the form of (2) may be inconsistent with the adjustment rules that are optimal given their priors (which in general may depend nonlinearly, and nonseparably, on the observable history). It might be possible to find repeated-game equilibria, perhaps with payoff perturbations as in McKelvey and Palfrey's (1992) analysis of experiments with the centipede game, that are statistically consistent with VHBB's subjects' behavior. But for this application it is both simpler and, given the unimportance of individual subjects' influences on y_t noted above, more plausible to adopt the working hypothesis that players focus on stage-game strategies.¹⁶ Under simple restrictions on the learning process this approach yields specific predictions that correspond to the dynamics VHBB observed. Although the possibility that players choose optimally given correct beliefs is not ruled out, it is strongly rejected in the empirical analysis.

To sum up, although the specification of players' learning rules used here cannot be defended by an appeal to rationality, it accommodates a wide range of sensible rules, which locate players' best replies quickly and reliably. The case for this specification is most compelling when it is difficult to defend specific restrictions on the form of players' Bayesian priors. However, the adjustments it

¹⁶ One might still try to impose rationality in the weaker sense of requiring players' rules to be statistically optimal for the process implied by the model, on the assumption that their influences on y_t are negligible. If players' choices are certainty-equivalent, (2) includes rules (exponential smoothing with time-varying coefficients) that are optimal in this sense for $\{y_t\}$ processes close to the one implied by the model (which is shown below to resemble a random walk plus noise, with drift and declining variances); see Harvey (1989). However, the logic of this approach also requires using the underlying variances of the process to determine the optimal coefficients of the learning rule. Even without the restrictive form of (2), attempts to close the loop in this way ultimately founder on the multiplicity of equilibria in the learning process, and in any case do not help to explain the differences in players' beliefs.

implies are qualitatively similar to those of Bayesian players and coincide with them in leading cases. Thus, using (2) might not be a significant source of error even if the population consisted entirely of Bayesians with known priors. More to the point, it appears to be a reasonable way to describe players' behavior when, as here, precise knowledge about their priors and decision rules is lacking.

When players' learning rules are described by (2), differences in their beliefs appear, to an outside observer or the players themselves, as random variations in the a_{it} and b_{it} . In keeping with their unpredictability, I characterize them statistically, treating certain parameters of their distributions as behavioral data. Thus, let

$$(3) \quad \varepsilon_{it} \equiv a_{it} - \alpha_t \quad \text{and} \quad \eta_{it} \equiv b_{it} - \beta_t \quad (t = 1, \dots),$$

where α_t and β_t are the common expected values of the a_{it} and b_{it} . The restrictions on the a_{it} and b_{it} discussed above suggest that $\alpha_t \rightarrow 0$ as $t \rightarrow \infty$ (but that the α_t may have either sign before they converge) and that $0 < \beta_t \leq 1$. By construction, $E\varepsilon_{it} = E\eta_{it} = 0$ for all i and t , where Ez denotes the expected value of the random variable z ex ante (that is, before random variables dated 0 are drawn). I assume that the ε_{it} and η_{it} are serially independent. This amounts to assuming that x_{it-1} fully captures any future effects of idiosyncratic influences on player i 's beliefs through period $t - 1$, which is restrictive, but seems a reasonable simplification. It is then a natural extension of my assumption that players are indistinguishable ex ante to assume that the ε_{it} and η_{it} are jointly independently and identically distributed (henceforth, "i.i.d.") across i for any given t , with exogenous common variances and covariances, denoted $\sigma_{\varepsilon t}^2$, $\sigma_{\eta t}^2$, and κ_t respectively.¹⁷

It is convenient to rewrite the model by substituting (3) into (1) and (2) and letting

$$(4) \quad \zeta_{i0} \equiv \varepsilon_{i0} \quad \text{and} \quad \zeta_{it} \equiv \varepsilon_{it} + (y_{t-1} - x_{it-1})\eta_{it} \quad (t = 1, \dots),$$

which yields

$$(5) \quad x_{i0} = \alpha_0 + \zeta_{i0}$$

and

$$(6) \quad x_{it} = \alpha_t + \beta_t y_{t-1} + (1 - \beta_t)x_{it-1} + \zeta_{it} \quad (t = 1, \dots).$$

The ζ_{it} are idiosyncratic random variables that represent the differences in players' initial beliefs and in how they respond to new observations. It is easily verified that, like the ε_{it} and η_{it} , they are ex ante identically distributed across i for any given t . Further, letting $E_s(z|\cdot)$ represent the conditional expectation of the random variable z given the information available at time s (that is, after random variables dated s are drawn) and using the law of iterated expectations,

¹⁷ Note that it is *not* assumed that ε_{it} and η_{it} are independent of each other. This would be inappropriate because a_{it} and b_{it} are different aspects of the same learning rule.

$E[\zeta_{it}\zeta_{jt}] = E[E_{t-1}(\zeta_{it}\zeta_{jt}|x_{t-1})] = 0$ for any $i \neq j$ and any t ; and $E\zeta_{it} = 0$, $E_s(\zeta_{it}|\cdot) = 0$, and $E[\zeta_{is}\zeta_{jt}] = E[E_s(\zeta_{is}\zeta_{jt}|x_s, \zeta_{is})] = E[\zeta_{is}E_s(\zeta_{jt}|x_s)] = 0$ for any i, j and any s, t with $s < t$. Thus, although the distribution of ζ_{it} depends on x_{t-1} and therefore on ζ_{it-1} and the other ζ_{jt-1} , the ζ_{it} are uncorrelated across i for any given t and serially uncorrelated. The common ex ante variance of the ζ_{it} , which is endogenous but well defined once $\sigma_{\epsilon t}^2$, $\sigma_{\eta t}^2$, and κ_t are specified, is denoted $\sigma_{\zeta t}^2$, and the conditional variance of ζ_{it} given x_{t-1} is denoted $\sigma_{\zeta it}^2|x_{t-1}$.

Under these distributional assumptions, (5) and (6) define a Markov process with state vector x_t , in which players' beliefs and strategy choices are identically distributed ex ante, but not in general otherwise. The recursive structure, together with the conditional independence of players' deviations from the average learning rule, captures the requirement that players must form their beliefs and choose their strategies independently which is the essence of the coordination problem.

The dynamics are driven by the dispersion of players' beliefs, as represented by the $\sigma_{\zeta it}^2|x_{t-1}$. This is true even though the model is formally consistent with any history of the y_t for any n and $f(\cdot)$, with the α_t varying as necessary over time and players' beliefs constrained to be identical throughout. (If $\sigma_{\zeta it}^2|x_{t-1} = 0$ for all i and t , so that $\zeta_{it} = 0$ with probability one, then (5) and (6) imply that $x_{it} = y_t = \sum_{s=0}^t \alpha_s$.) Solutions of this kind, in which players jump simultaneously from one stage-game equilibrium to the next following some commonly understood pattern, correspond to equilibria of the repeated game that are difficult to rule out using traditional refinements. Yet it is clear that they provide neither a meaningful explanation of the dynamics of y_t nor an adequate statistical summary of subjects' strategy choices. I shall therefore ask that the model meet both of these goals simultaneously, without such ad hoc variations in its parameters.

This requires that players differ significantly in their responses to new information as well as in their initial beliefs. To see this, assume for simplicity that $\alpha_t = 0$ for all $t = 1, \dots$. If $\sigma_{\zeta 0}^2 = \sigma_{\epsilon 0}^2 > 0$ but $\sigma_{\zeta it}^2|x_{t-1} = 0$ for all i and $t = 1, \dots$, then (6) (with $\zeta_{it} = \alpha_t = 0$ and $0 < \beta_t \leq 1$) implies that $(x_{it} - y_{t-1})$ and $(x_{it-1} - y_{t-1})$ always have the same sign, with x_{it} closer to y_{t-1} than x_{it-1} was. It follows that players' strategy choices converge to y_t monotonically, without overshooting, and therefore, when $f(\cdot)$ is an order statistic, that $y_t \equiv y_0$ for all t , independent of n and $f(\cdot)$. This extreme history-dependence is consistent with the results for VHBB's median treatments but not for their minimum treatments.¹⁸

Given that $\sigma_{\zeta it}^2|x_{t-1} > 0$, the model is closest to other work on adaptive dynamics when $\alpha_t = 0$ for all $t = 1, \dots$ and β_t and $\sigma_{\zeta it}^2|x_{t-1}$ remain constant at

¹⁸ In the large-group minimum treatment, for instance, subjects whose efforts were above the minimum adjusted their efforts only part of the way toward it on average. Thus, without persistent differences in subjects' responses to new information the minimum would never have fallen. In fact, however, there was enough variation in subjects' responses, given the size of the population, to make the minimum fall in 9 out of the 13 instances in which it was not already at the lowest possible level.

positive levels or converge to positive limits, so that the process has stationary or asymptotically stationary transition probabilities and a unique, globally stable ergodic distribution. Then, the methods of Freidlin and Wentzell (1984), as first applied in this area by Foster and Young (1990) and subsequently adapted to discrete state spaces by Kandori, Mailath, and Rob (1993) and Young (1993a), can be applied to the model with discrete strategy choice analyzed in Section 5. In the long run the process cycles perpetually among the pure-strategy equilibria of the underlying game, with their prior probabilities at any given time determined by the ergodic distribution. As the $\sigma_{\epsilon_{it}}^2|x_{t-1}$ are allowed to approach zero (remaining constant over time) the ergodic distribution assigns probability approaching one to the equilibrium with the lowest effort whenever the summary statistic is below the median, and approaches a limit with positive probability on every equilibrium when the summary statistic is the median, in each case independent of the number of players.¹⁹

Thus, an analysis in the style of Kandori, Mailath, and Rob (1993) and Young (1993a) discriminates between VHBB's large-group minimum treatment on one hand and their median treatments and random-pairing minimum treatment (in which the relevant summary statistic is also the median, as explained in footnote 10) on the other, but otherwise does not distinguish between them. These conclusions are not logically inconsistent with the systematic, persistent effects of the number of players and the summary statistic within the finite horizons of VHBB's experiments, but they suggest that an analysis based entirely on the limiting distributions of an ergodic learning process with infrequent "mutations" is of limited use in understanding those effects.

The methods developed below characterize the probability distributions of the entire time paths of y_t and the x_{it} , whether or not the process is ergodic and the $\sigma_{\epsilon_{it}}^2|x_{t-1}$ are small. As will become clear, the model's dynamics are closest to VHBB's results when the $\sigma_{\epsilon_{it}}^2|x_{t-1}$ decline steadily over time—even-
tually to zero, given subjects' tendency to "lock in" on an equilibrium—just as one would expect as players gain experience forecasting y_t .

4. ANALYSIS

The model's "evolutionary" structure makes it analytically tractable despite its nonlinearity and the nonstationarity of its transition probabilities. When players' interactions are filtered through a summary statistic $f(x_{1t}, \dots, x_{nt})$ with the property that $f(a + bx_{1t}, \dots, a + bx_{nt}) \equiv a + bf(x_{1t}, \dots, x_{nt})$, the outcome can be expressed as a function of the idiosyncratic shocks that represent the differences between their beliefs. In what follows, sums with no terms (like $\sum_{s=0}^{t-1} \beta_{s+1} f_s$ for $t = 0$) should be understood to equal 0, and products with no terms should be understood to equal 1. All limits are taken as $t \rightarrow \infty$.

¹⁹ This follows because fewer changes in individual players' efforts are required to get from a high-effort equilibrium to the basin of attraction of a low-effort equilibrium than vice versa if, and only if, the summary statistic is below the median. Robles (1994) analyzes a closely related model for the games studied here, showing that this conclusion changes only slightly when players do not ignore their influences on y_t .

PROPOSITION 1: *The unique solution of (5) and (6) is given, for all i and t , by*

$$(7) \quad x_{it} = \sum_{s=0}^t \alpha_s + \sum_{s=0}^{t-1} \beta_{s+1} f_s + z_{it}$$

and

$$(8) \quad y_t = \sum_{s=0}^t \alpha_s + \sum_{s=0}^{t-1} \beta_{s+1} f_s + f_t,$$

where

$$(9) \quad z_{it} \equiv \sum_{s=0}^t \left[\prod_{j=1}^{t-s} (1 - \beta_{t-j+1}) \right] \zeta_{is} \quad \text{and} \quad f_t \equiv f(z_{1t}, \dots, z_{nt}).$$

PROOF: The solution follows immediately by induction on t , recalling that $f(a + bx_{1t}, \dots, a + bx_{nt}) \equiv a + bf(x_{1t}, \dots, x_{nt})$ and noting that $z_{it} = (1 - \beta_t)z_{it-1} + \zeta_{it}$ for all t . Uniqueness also follows immediately by induction. Q.E.D.

Proposition 1 represents the outcome of the learning dynamics as the cumulative effect of trend and shock terms from each period, each of whose effects persist indefinitely. (The remaining terms, z_{it} in (7) and f_t in (8), are subsumed in the shock terms after the period in which they first appear.) Thus the coordination process resembles a random walk, but with declining variances and possibly nonzero drift. Although Proposition 1's solution remains valid no matter how the ζ_{it} are generated, much of its usefulness in the analysis below stems from the fact that the shock terms are known functions of the z_{it} , which under my assumptions are *ex ante* identically distributed and uncorrelated across i for all t .

The next proposition provides conditions under which the dynamics converge, with probability 1, to one of the symmetric Nash equilibria of the underlying coordination game. In this proposition, and sometimes below, it is necessary to bound players' strategies. This is accomplished by increasing x_{it} to its lower bound, denoted \underline{x} , or reducing it to its upper bound, denoted \bar{x} , whenever it would otherwise fall outside the interval $[\underline{x}, \bar{x}]$. This is equivalent to increasing ζ_{i0} to $\underline{x} - \alpha_0$ (or reducing it to $\bar{x} - \alpha_0$) when it would otherwise fall below (above) that value and increasing ζ_{it} to $\underline{x} - \alpha_0 - \beta_t y_{t-1} - (1 - \beta_t)x_{it-1}$ (or reducing it to $\bar{x} - \alpha_0 - \beta_t y_{t-1} - (1 - \beta_t)x_{it-1}$) when it would otherwise fall below (above) that value. The resulting truncation of the conditional distributions of the ζ_{it} is consistent with the distributional assumptions introduced in Section 3, except that it may bias the $E\zeta_{i0}$ or the $E_{t-1}(\zeta_{it}|x_{t-1})$ away from zero and induce serial or contemporaneous correlation in the ζ_{it} ; I indicate below when this affects the results.

PROPOSITION 2: *Assume that the distributions of the ζ_{it} are truncated so that the x_{it} always remain in the interval $[\underline{x}, \bar{x}]$. Then if β_t is bounded above 0, with*

$0 < \underline{\beta} \leq \beta_t \leq 1$ for all t , and $\sum_{s=0}^{\infty} \alpha_s$ and $\sum_{s=0}^{\infty} \sigma_{\xi_s}^2$ are finite, y_t and the x_{it} converge, with probability 1, to a common, finite limit.

PROOF: The proof follows the martingale convergence arguments of Nevel'son and Has'minskii (1973, Theorem 2.7.3).²⁰ Consider the Lyapunov function $V_t \equiv \sum_{i,j} (x_{it} - x_{jt})^2$, where the summation is taken over all $i, j = 1, \dots, n$. Clearly, $V_t \geq 0$ for all x_t , with $V_t = 0$ if and only if $x_{it} = x_{jt}$ for all i and j . Substituting from (6) and simplifying,

$$\begin{aligned} (10) \quad V_t &= \sum_{i,j} [(1 - \beta_t)(x_{it-1} - x_{jt-1}) + \zeta_{it} - \zeta_{jt}]^2 \\ &= \sum_{i,j} [(1 - \beta_t)^2 (x_{it-1} - x_{jt-1})^2 \\ &\quad + 2(1 - \beta_t)(x_{it-1} - x_{jt-1})(\zeta_{it} - \zeta_{jt}) + (\zeta_{it} - \zeta_{jt})^2]. \end{aligned}$$

Taking expectations in (10) yields

$$\begin{aligned} (11) \quad E_{t-1}(V_t | x_{t-1}) &= (1 - \beta_t)^2 \sum_{i,j} (x_{it-1} - x_{jt-1})^2 \\ &\quad + 2(1 - \beta_t) \sum_{i,j} (x_{it-1} - x_{jt-1}) E_{t-1}[(\zeta_{it} - \zeta_{jt}) | x_{t-1}] \\ &\quad + \sum_{i,j} E_{t-1}[(\zeta_{it} - \zeta_{jt})^2 | x_{t-1}]. \end{aligned}$$

The first term on the right-hand side of (11) is plainly bounded below V_{t-1} for all x_{t-1} outside any given neighborhood of the set for which $V_{t-1} = 0$. Without truncation the second term equals 0, and the third term eventually approaches 0 with probability 1 because $\sigma_{\xi_t}^2$ must approach 0 for $\sum_{s=0}^{\infty} \sigma_{\xi_s}^2$ to be finite and this cannot occur unless $E_{t-1}(\xi_{it}^2 | x_{t-1})$ approaches 0 with probability 1. Thus, without truncation $\{V_t\}$ eventually becomes a nonnegative supermartingale, and therefore converges with probability 1, under the stated variance condition, to its lower bound, 0. Because truncation at time t can never increase $E_{t-1}(V_t | x_{t-1})$ or $E_{t-1}(\xi_{it}^2 | x_{t-1})$, this conclusion extends to the truncated version of the $\{V_t\}$ process. It follows that $(x_{it} - x_{jt})$ converges to 0, with probability 1, for all $i, j = 1, \dots, n$, and therefore (because $y_t \equiv f(x_t)$, $f(\cdot)$ is continuous, and $f(a, \dots, a) \equiv a$) that $(y_t - x_{it})$ also converges to 0 with probability 1. The convergence of the x_{it} (and not just their differences) then follows, given that they are bounded, from two observations: (i) x_t cannot (with positive probability) return infinitely often to a given point at which $x_{it} \neq y_t$ for some i , because

²⁰ Nevel'son and Has'minskii's (1973) generalization of Lyapunov's global-stability criterion to stochastic dynamic systems is a useful alternative to the techniques of Ljung (see Ljung and Soderstrom (1983)) used by Woodford (1990). Nevel'son and Has'minskii require that the Lyapunov function approach infinity with $\|x_t\|$, but they use this condition only to ensure the boundedness of solution paths; in Proposition 2 boundedness is assumed directly. The variance condition they impose is just what one would expect by analogy with the strong law of large numbers.

the $(y_t - x_{it})$ converge to 0 with probability 1; and (ii) x_t cannot (with positive probability) oscillate infinitely often between distinct points at which $x_{it} = y_t$ for all i , because $\sum_{s=0}^t \alpha_s$ and the $\sum_{s=0}^t \zeta_{is}$ also converge (the latter with probability 1) under the stated conditions. *Q.E.D.*

Because the x_{it} , and therefore the $(x_{it} - y_t)$, are bounded, (4) implies that for $\sum_{s=0}^{\infty} \sigma_{\zeta_s}^2$ to be finite it is sufficient that $\sum_{s=0}^{\infty} \sigma_{\varepsilon_s}^2$ and $\sum_{s=0}^{\infty} \sigma_{\eta_s}^2$ be finite, and necessary that $\sum_{s=0}^{\infty} \sigma_{\varepsilon_s}^2$ be finite. Although it is plainly not necessary that $\sum_{s=0}^{\infty} \sigma_{\eta_s}^2$ be finite, I have not been able to use this possibility to weaken these sufficient conditions for convergence in any useful way. I have also not been able, without truncating the distributions of the ζ_{it} , to rule out the possibility of solutions in which y_t and the x_{it} diverge together to $\pm\infty$, but the conclusion of Proposition 2 seems likely to remain valid for well-behaved unbounded distributions.

Propositions 1 and 2 show that the model determines a prior probability distribution over limiting equilibria that is normally nondegenerate, rather than singling out a particular equilibrium. This is consistent with the variations observed across different runs in some of VHBB's treatments, and follows naturally from characterizing players' beliefs statistically instead of assuming that they can be determined entirely within the model.

I now consider the model's comparative dynamics properties, focusing on how changes in VHBB's treatment variables, the number of players, n , and the summary statistic, $f(\cdot)$, affect the prior probability distributions of the y_t and x_{it} .²¹ Propositions 3 and 4 characterize these effects qualitatively for given values of the behavioral parameters α_t , β_t , $\sigma_{\varepsilon_t}^2$, $\sigma_{\eta_t}^2$, and κ_t , showing that they are fully consistent with the patterns of variation in the dynamics VHBB observed across treatments. Propositions 5 and 6 then show how the means and variances of the y_t and x_{it} are determined by the treatment variables and the parameters, to assess the quantitative magnitudes of these effects.

These results are then used to explain the patterns of variation in the dynamics. As with any model whose predictions depend on empirical parameters, the explanation rests on the assumption that the parameters are stable across runs in any given treatment. (Sample-splitting tests and inspecting the data suggest that this is not a bad assumption.) But something more is involved here, because each treatment is a different game, and there is no good theoretical reason to expect players' responses to be the same in different games. Thus, in principle, the effects of changes in the treatment variables identified in Propositions 3 and 4 might be swamped by large, unpredictable changes in the behavioral parameters. Such changes would leave the model's ability to predict coordination outcomes unaffected, as long as the behavioral

²¹ Broseta (1993a) studies the comparative dynamics of a closely related model that allows for autoregressive conditional heteroskedasticity in the idiosyncratic components of players' beliefs, obtaining some results like those given here and others concerning the effects of changes in the initial mean and dispersion of players' beliefs.

parameters remained stable across runs in each treatment. They would, however, reduce the patterns of variation in the dynamics to a primarily empirical question, limiting the theory's usefulness beyond those environments for which the parameters have been estimated and leaving it unable to explain why the dynamics varied across treatments as they did in VHBB's experiments. There would then be no reason to expect similar results or patterns of variation in the dynamics for nearby values of the treatment variables.

This question is addressed below by showing that the most important differences VHBB observed across treatments were mainly the product of the interactions between the number of players, the summary statistic, and strategic uncertainty characterized in Propositions 3–6. Unreported simulations of the present model and the sensitivity analysis of a closely related model in Broseta (1993a) reveal that the patterns in the dynamics are highly robust to moderate changes in the behavioral parameters, which are shown in the econometric analysis in Section 6 to vary surprisingly little across treatments. Estimates based on Propositions 5 and 6 suggest that the effects of the important changes in treatment variables were systematic and large enough to swamp the effects of the changes in the behavioral parameters they induced; this is confirmed in Section 7 when the model is simulated using Section 6's estimates. Thus, there is good reason to expect similar results in nearby environments, and to believe that the patterns VHBB observed are replicable.

In what follows, the j th order statistic of an empirical distribution (ξ_1, \dots, ξ_n) is defined to be the j th 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, in general, to a shift in the weights used to compute $f(\cdot)$ that increases j in the sense of first-order stochastic dominance, viewing the weights as a probability distribution for the purpose of applying the definition. A random variable is said to "stochastically increase" if its probability distribution shifts upward in the sense of first-order stochastic dominance.

PROPOSITION 3: *Suppose that $b_{it} \geq 0$ for all i and t , so that $\partial x_{it}/\partial y_{t-1} \geq 0$ with probability one. Then, holding n constant, increasing j stochastically increases y_t and the x_{it} for any t .*

PROOF: The proof is immediate from (2) by induction, noting that for any given value of x_{t-1} , y_{t-1} , and therefore x_t , is weakly larger when computed for the larger value of j . Q.E.D.

PROPOSITION 4: *Suppose that $b_{it} \geq 0$ for all i and t , so that $\partial x_{it}/\partial y_{t-1} \geq 0$ with probability one. Then, holding j (or the weights on alternative values of j) constant, increasing n stochastically decreases y_t and the x_{it} for any t .*

PROOF: The proof is immediate from (2) by induction, noting that for any given value of x_{t-1} (with i running from 1 to the larger value of n), y_{t-1} , and therefore x_t , is weakly smaller when computed for the larger value of n . *Q.E.D.*

Propositions 3 and 4 hold whether or not the distributions of the ζ_{it} are truncated. They make precise, in the probabilistic sense appropriate to the model, the common intuitions that coordination is more difficult, the smaller the subsets of the population that can adversely affect the outcome, and that coordination is more difficult in larger groups because it requires coherence among a larger number of independent decisions.

Let $\sigma_{z_t}^2$ denote the common ex ante variance of the z_{it} . Because the ζ_{it} are serially uncorrelated, it follows from (9) that $\sigma_{z_t}^2 \equiv \sum_{s=0}^t [\prod_{j=1}^{t-s} (1 - \beta_{t-j+1})]^2 \sigma_{\zeta_s}^2$. Define $\mu_t \equiv Ef(z_{1t}/\sigma_{z_t}, \dots, z_{nt}/\sigma_{z_t})$. Because the random variables z_{it}/σ_{z_t} are standardized, with common mean 0 and common variance 1, μ_t is completely determined by n , $f(\cdot)$, and the joint distribution of the z_{it}/σ_{z_t} . μ_t is subscripted only because the distribution of the z_{it}/σ_{z_t} is generally time-dependent; its dependence on n and $f(\cdot)$ is suppressed for clarity. It is easily shown, using the properties of order statistics, that μ_t is decreasing in n and increasing in j and that, if $n \geq 2$ and the z_{it} are distributed symmetrically about 0, then $\mu_t < 0$ if $1 \leq j < (n+1)/2$, $\mu_t = 0$ if $j = (n+1)/2$, and $\mu_t > 0$ if $(n+1)/2 < j \leq n$.

PROPOSITION 5: *The ex ante means of y_t and the x_{it} are given, for all i and t , by*

$$(12) \quad Ex_{it} = \sum_{s=0}^t \alpha_s + \sum_{s=0}^{t-1} \beta_{s+1} \sigma_{zs} \mu_s$$

and

$$(13) \quad Ey_t = \sum_{s=0}^t \alpha_s + \sum_{s=0}^{t-1} \beta_{s+1} \sigma_{zs} \mu_s + \sigma_{zt} \mu_t.$$

PROOF: Taking expectations in (7) and (8), using (9), and noting that

$$(14) \quad Ef(z_{1s}, \dots, z_{ns}) \equiv E[\sigma_{zs} f(z_{1s}/\sigma_{zs}, \dots, z_{ns}/\sigma_{zs})] \equiv \sigma_{zs} \mu_s$$

immediately yields (12) and (13). *Q.E.D.*

Let $\text{Var } z$ and $\text{Cov}(z, z')$ denote the ex ante variance of the random variable z and the covariance of the random variables z and z' . Let $\sigma_t^2 \equiv \text{Var } f(z_{1t}/\sigma_{z_t}, \dots, z_{nt}/\sigma_{z_t})$, $\gamma_t \equiv E[(z_{it}/\sigma_{z_t}) f(z_{1t}/\sigma_{z_t}, \dots, z_{nt}/\sigma_{z_t})]$ (for any i), and $\delta_{st} \equiv \text{Cov}[f(z_{1s}/\sigma_{zs}, \dots, z_{ns}/\sigma_{zs}), f(z_{1t}/\sigma_{z_t}, \dots, z_{nt}/\sigma_{z_t})]$. Like μ_t , the parameters σ_t^2 , γ_t , and δ_{st} are completely determined by n , $f(\cdot)$, and the joint distribution of the z_{is}/σ_{zs} and z_{it}/σ_{z_t} ; they are subscripted only because that distribution is generally time-dependent.

PROPOSITION 6:²² *The ex ante variances of y_t and the x_{it} are given, for all i and t , by*

$$(15) \quad \text{Var } x_{it} = \sum_{s=0}^{t-1} \beta_{s+1}^2 \sigma_{zs}^2 \sigma_s^2 + \sum_{s=0}^{t-1} \beta_{s+1} \sigma_{zs}^2 \prod_{j=1}^{t-s} (1 - \beta_{t-j+1}) \gamma_s \\ + 2 \sum_{s=1}^{t-1} \sum_{r=0}^{s-1} \beta_{r+1} \beta_{s+1} \sigma_{zr} \sigma_{zs} \delta_{rs} + \sigma_{zt}^2$$

and

$$(16) \quad \text{Var } y_t = \sum_{s=0}^{t-1} \beta_{s+1}^2 \sigma_{zs}^2 \sigma_s^2 + \sum_{s=0}^{t-1} \beta_{s+1} \sigma_{zs} \sigma_{zt} \delta_{st} \\ + 2 \sum_{s=1}^{t-1} \sum_{r=0}^{s-1} \beta_{r+1} \beta_{s+1} \sigma_{zr} \sigma_{zs} \delta_{rs} + \sigma_{zt}^2 \sigma_t^2.$$

PROOF: Taking variances in (7) and (8) yields

$$(17) \quad \text{Var } x_{it} = \sum_{s=0}^{t-1} \beta_{s+1}^2 \text{Var } f_s + \sum_{s=0}^{t-1} \beta_{s+1} \text{Cov}(f_s, z_{it}) \\ + 2 \sum_{s=1}^{t-1} \sum_{r=0}^{s-1} \beta_{r+1} \beta_{s+1} \text{Cov}(f_r, f_s) + \text{Var } z_{it}$$

and

$$(18) \quad \text{Var } y_t = \sum_{s=0}^{t-1} \beta_{s+1}^2 \text{Var } f_s + \sum_{s=0}^{t-1} \beta_{s+1} \text{Cov}(f_s, f_t) \\ + 2 \sum_{s=1}^{t-1} \sum_{r=0}^{s-1} \beta_{r+1} \beta_{s+1} \text{Cov}(f_r, f_s) + \text{Var } f_t.$$

Because the z_{is}/σ_{zs} are standardized, with mean 0 and variance 1,

$$(19) \quad \text{Var } f_s \equiv \text{Var } f(z_{1s}, \dots, z_{ns}) \\ \equiv \text{Var} [\sigma_{zs} f(z_{1s}/\sigma_{zs}, \dots, z_{ns}/\sigma_{zs})] \equiv \sigma_{zs}^2 \sigma_s^2.$$

For any $t > s$, z_{it} in (9) can be expressed as

$$(20) \quad z_{it} = \prod_{j=1}^{t-s} (1 - \beta_{t-j+1}) z_{is} + \sum_{k=1}^{t-s} \left[\prod_{j=1}^{t-s-k} (1 - \beta_{t-j+1}) \right] \zeta_{is+k}.$$

Because $Ez_{it} = 0$ and $E(z_{is}\zeta_{it}) = E[E_s(z_{is}\zeta_{it}|x_s, z_{is})] = E[z_{is}E_s(\zeta_{it}|x_s)] = 0$ for

²² Because the proof of Proposition 6 uses the properties of the untruncated distributions of the ζ_{it} , its formulas would need to be modified with truncation.

any $s < t$, it follows from (20) and my assumptions on $f(\cdot)$ that, for all $s < t$,

$$\begin{aligned}
 (21) \quad \text{Cov}(f_s, z_{it}) &= E[z_{it} f(z_{1s}, \dots, z_{ns})] \\
 &= \sigma_{zs}^2 E[(z_{is}/\sigma_{zs}) f(z_{1s}/\sigma_{zs}, \dots, z_{ns}/\sigma_{zs})] \\
 &\quad \times \prod_{j=1}^{t-s} (1 - \beta_{t-j+1}) \\
 &= \sigma_{zs}^2 \gamma_s \prod_{j=1}^{t-s} (1 - \beta_{t-j+1})
 \end{aligned}$$

and

$$\begin{aligned}
 (22) \quad \text{Cov}(f_s, f_t) &\equiv \text{Cov}[f(z_{1s}, \dots, z_{ns}), f(z_{1t}, \dots, z_{nt})] \\
 &\equiv \sigma_{zs} \sigma_{zt} \text{Cov}[f(z_{1s}/\sigma_{zs}, \dots, z_{ns}/\sigma_{zs}), \\
 &\quad f(z_{1t}/\sigma_{zt}, \dots, z_{nt}/\sigma_{zt})] \\
 &\equiv \sigma_{zs} \sigma_{zt} \delta_{st}.
 \end{aligned}$$

Substituting (19), (21), and (22) into (17) and (18) yields (15) and (16). *Q.E.D.*

Techniques like those used to prove Propositions 5 and 6 make it possible to compute the ex ante variances $\sigma_{\xi t}^2$ and $\sigma_{z t}^2$ recursively from the fundamental parameters. Given (4) and (9),

$$(23) \quad \sigma_{z0}^2 = \sigma_{\xi0}^2 = \sigma_{\varepsilon0}^2$$

and

$$(24) \quad \sigma_{z t}^2 = (1 - \beta_t)^2 \sigma_{z t-1}^2 + \sigma_{\xi t}^2.$$

Combining (4), (7), (8), (14), and (21) and recalling that $\kappa_t \equiv \text{Cov}(\varepsilon_{it}, \eta_{it})$ yields

$$(25) \quad \sigma_{\xi t}^2 = \sigma_{\varepsilon t}^2 + 2\sigma_{z t-1} \mu_t \kappa_t + \sigma_{z t-1}^2 (\sigma_t^2 + 1 + \mu_t^2 - 2\gamma_t) \sigma_{\eta t}^2.$$

Intuition suggests that both $\sigma_{\xi t}^2$ and $\sigma_{z t}^2$ should decline over time as players learn to forecast y_t . It is clear from (23) and (24) that (with $0 < \beta_t \leq 1$) the $\sigma_{\xi t}^2$ are declining whenever the $\sigma_{z t}^2$ are, so that requiring the $\sigma_{z t}^2$ to decline is more stringent than requiring the $\sigma_{\xi t}^2$ to decline.

The formulas in Propositions 5 and 6 show with considerable generality how the dispersion of players' beliefs interacts with the strategic environment to determine the mean and variance of the outcome. But before those formulas can be put to practical use, numerical values must be assigned to the parameters that appear in them. As noted above, the behavioral parameters α_t , β_t , and $\sigma_{\varepsilon t}^2$, $\sigma_{\eta t}^2$, and κ_t (which determine the $\sigma_{z t}^2$) must be estimated for each environment; this is done for VHBB's experiments in Section 6. The remaining parameters μ_t , σ_t^2 , γ_t , and δ_{st} are completely determined by n , $f(\cdot)$, and the distributions of the z_{it} (which are determined by the distributions of the ε_{it} and η_{it}), but are difficult to evaluate due to the complexity of their dependence on those distributions.

This problem can be overcome in two ways. Evaluating μ_t , σ_t^2 , γ_t , and δ_{st} can be sidestepped by estimating the probability distributions of outcomes implied

by the model directly by repeated simulation; this is done in Section 7, using the parameter estimates reported in Section 6. Alternatively, the outcome distributions can be approximated analytically under simplifying assumptions about the distributions of the z_{it} .

Because the analytical results provide insight into the workings of the model, I outline the approximation method here.²³ The approximations are based on the assumption that the z_{it} are jointly normally distributed for any given t . Normality is a reasonable approximation, at least when truncation effects are negligible, because the z_{it} are weighted sums of the ζ_{it} , which are weakly dependent and likely to be approximately conditionally normal for familiar reasons.²⁴ Normality simplifies matters by making the common distribution of the z_{it}/σ_{zt} independent of t , so that $\mu_t \equiv \mu$, $\sigma_t^2 \equiv \sigma^2$, and $\gamma_t \equiv \gamma$. Given that the z_{it}/σ_{zt} are uncorrelated for any given t , the parameters μ and σ^2 are tabulated in Teichroew (1956) for any order statistic of the normal distribution and any $n \leq 20$; and it can be shown that $\gamma = 1/n$.²⁵ Finally, for jointly normal z_{is} and z_{it} , δ_{st} is completely determined by the common correlation between the standardized z_{is}/σ_{zs} and the corresponding z_{it}/σ_{zt} . Because $E(z_{is}\zeta_{it}) = 0$ whenever $s < t$, (20) implies that this correlation then equals

$$(26) \quad \text{Cov}(z_{is}, z_{it})/\sigma_{zs}\sigma_{zt} = \prod_{j=1}^{t-s} (1 - \beta_{t-j+1}) (\text{Var } z_{is})/\sigma_{zs}\sigma_{zt} \\ = \prod_{j=1}^{t-s} (1 - \beta_{t-j+1}) \sigma_{zs}/\sigma_{zt}.$$

²³ An earlier version, Crawford (1992), provides more detail about the approximations and compares them with the probabilities estimated directly by repeated simulation.

²⁴ It does not seem possible to give an exact justification for this assumption. Normality of the z_{it} would follow from joint normality of the ζ_{it} , even though they are dependent, because the normal distributions form a stable class. But even if the ε_{it} and η_{it} were normal the ζ_{it} could not be jointly normal, because then their lack of serial correlation would imply that they were serially independent, a contradiction.

²⁵ Let $f^j(\cdot)$ denote the j th order statistic. If ξ_1, \dots, ξ_n are any i.i.d. random variables, it follows from symmetry, the linearity of the expectations operator, and the fact that the sum of an empirical distribution's elements equals the sum of its order statistics that

$$E[f(\xi_1, \dots, \xi_n)\xi_i] = \frac{1}{n} \sum_{j=1}^n E[f(\xi_1, \dots, \xi_n)\xi_j] = \frac{1}{n} E \left[f(\xi_1, \dots, \xi_n) \sum_{j=1}^n \xi_j \right] \\ = \frac{1}{n} E \left[f(\xi_1, \dots, \xi_n) \sum_{j=1}^n f^j(\xi_1, \dots, \xi_n) \right] \\ = \frac{1}{n} \sum_{j=1}^n E[f(\xi_1, \dots, \xi_n) f^j(\xi_1, \dots, \xi_n)].$$

It is well known (see for example Jones (1948)) that when the ξ_i are standard normal, the sum in the last term equals 1 for any order statistic $f(\cdot)$; this conclusion extends immediately to the case where $f(\cdot)$ is a convex combination of order statistics.

Thus, under normality δ_{st} can be written as $\delta(\rho_{st})$, where $\rho_{st} \equiv \prod_{j=1}^{t-s} (1 - \beta_{t-j+1}) \sigma_{zs} / \sigma_{zt}$. It is clear that $0 \leq \rho_{st} \leq 1$, and that the function $\delta(\cdot)$ is differentiable and monotonically increasing on the interval $[0, 1]$, with $\delta(0) = 0$ and $\delta(1) = \sigma^2$. The function $\delta(\cdot)$ has not been tabulated, to my knowledge, but could be tabulated with some effort.

An explanation for the patterns of variation in the dynamics can now be discerned. (This explanation is made precise in Section 7, taking the discreteness of players' strategy choices into account.) For simplicity, ignore the discreteness and boundedness of y_t and the x_{it} , assume normality of the z_{it} , and suppose that $\alpha_s \approx 0$ and $\beta_s \approx \beta$ for all $s = 1, \dots$ and that $\sum_{s=0}^t \sigma_{zs} \rightarrow S$. Proposition 5 then implies that Ey_t and Ex_{it} approach the approximate common limit $\alpha_0 + \mu\beta S$. Given that μ is determined by n and $f(\cdot)$, this formula shows how the means and dispersions of players' beliefs and the average rate at which they respond to new information interact with the number of players and the summary statistic to determine the ex ante mean of the limiting coordination outcome.

For VHBB's random-pairing minimum treatment (viewed as a median treatment, as in footnote 10) and their median treatments, $\mu = 0$ because of symmetry. Thus, in these treatments the approximate common limit of Ey_t and Ex_{it} is α_0 . The estimates of α_0 in Section 6 range from 4.30 in the random-pairing minimum treatment to 4.71 and 4.75 in median treatments Φ and Γ and 6.26 in median treatment Ω (whose payoff structure, described in footnote 34 below, made the efficient equilibrium with effort 7 more prominent than in the other treatments). For VHBB's large-group minimum treatment, $\mu = -1.74$.²⁶ The estimates of α_0 and β are 5.45 and 0.25 respectively, and S (which is difficult to estimate for this treatment, due to boundary effects) appears highly unlikely to be less than 10. Thus, the approximate common limit of Ey_t and Ex_{it} , $\alpha_0 + \mu\beta S$, is at most 1.10.

Even the simplifying assumptions made above in approximating the means do not yield simple expressions for $\text{Var } y_t$ and $\text{Var } x_{it}$, but Proposition 6 can be used to get some idea of how they vary with the environment. Because $\delta(\rho_{st})$ increases smoothly from 0 to σ^2 as ρ_{st} ranges from 0 to 1, its value can be roughly approximated by $\rho_{st}\sigma^2$. It then follows from (15) and (16) that $\text{Var } y_t$ and $\text{Var } x_{it}$ are approximately proportional to σ^2 for each treatment, with factors of proportionality determined by n , β , and the σ_{zt}^2 .

Like μ , σ^2 is determined by n and $f(\cdot)$. $\sigma^2 = 0.17$ in the median treatments, 0.10 in the random-pairing minimum treatment, and 0.30 in the large-group minimum treatment. The estimates of $\sigma_{\epsilon t}^2$ and $\sigma_{\eta t}^2$ reported in Section 6 are small for the median treatments, significantly larger for the large-group minimum treatment, and much larger for the random-pairing minimum treatment. This suggests that the ex ante variance of the limiting outcome is small in the

²⁶ This value for μ and the values for σ^2 given below are taken from Teichroew's (1956) Tables I and II, setting $n = 15$ in the large-group minimum treatment for simplicity.

median treatments, larger in the large-group minimum treatment, and still larger in the random-pairing minimum treatment.

These approximations suggest that the limiting prior probability distribution of y_i is relatively concentrated, with means 4.71, 4.75, and 6.26 in median treatments Φ , Γ , and Ω respectively; somewhat more diffuse, with mean no greater than 1.10, in the large-group minimum treatment; and still more diffuse, with mean 4.30, in the random-pairing minimum treatment. These estimated means are very close to those implied by the precise estimates of the limiting prior probability distributions implied by the model, and almost as close to those implied by the last-period frequency distributions in VHBB's experiments (both of these distributions are reported in Tables VI–Xa). The suggested ordering of the variances is also qualitatively consistent with the precisely estimated limiting distributions and VHBB's limiting frequency distributions, once the effect of the lower boundary in the large-group minimum treatment is taken into account. I do not estimate the variances here because this adds little intuition, and the distributions are estimated precisely below.

The most important changes in the dynamics across treatments VHBB observed were between the random-pairing and large-group minimum treatments, and between the median treatments and the large-group minimum treatment. Viewing the random-pairing minimum treatment as a median treatment, as explained in footnotes 10 and 11, the model treats the differences between these treatments primarily as changes in the summary statistic (even though the former difference is “really” a change in group size and the latter also involves a change in the size of the group, from 9 to about 15). The estimates of the limiting means of y_i given above suggest that each of these changes altered the drift of the process by much more than the accompanying changes in the behavioral parameters.

VHBB also observed differences in the dynamics between the median treatments and the random-pairing minimum treatment. These were generally smaller and of a different character, having to do with the rate of convergence and, in one case, the mean outcome. The model treats the differences between these treatments mainly as changes in group size, which are relatively small and in any case have no effect on drift. The above estimates of the limiting means of y_i suggest that the effects of these changes in treatment variables are mainly determined by changes in the behavioral parameters.

5. DISCRETE STRATEGY CHOICE

I now extend the analysis to the discrete strategy spaces of VHBB's experimental environments. The x_{it} will continue to represent players' beliefs, and remain continuously variable, but instead of representing players' strategy choices directly they will now determine them as the latent variables in a discrete-choice model. The fact that players' strategy choices and best replies are naturally ordered by their payoff implications suggests modeling them as an ordered probit (see for example McFadden (1984)). I therefore assume that the

ζ_{it} are conditionally normally distributed given x_{t-1} , and that players' choices among the feasible discrete alternatives are determined by rounding the latent variables that represent their beliefs as described below. These assumptions determine players' choice probabilities as functions of the nonrandom components of their latent variables.

Let players choose, as in VHBB's experiments, among seven pure strategies, numbered $1, \dots, 7$. I assume that player i 's choice at time t is determined by rounding x_{it} to the nearest feasible strategy, so that he chooses strategy 1 when $x_{it} \leq 3/2$; strategy x , for $x = 2, \dots, 6$, when $x - 1/2 < x_{it} \leq x + 1/2$; and strategy 7 when $13/2 < x_{it}$.²⁷ Conditional on x_{t-1} , ζ_{it} is normal with mean 0 and variance (computed from (4))

$$(27) \quad \sigma_{\zeta_{it}}^2 | x_{t-1} = \sigma_{\epsilon_t}^2 + 2(y_{t-1} - x_{it-1})\kappa_t + (y_{t-1} - x_{it-1})^2 \sigma_{\eta_t}^2.$$

Thus, letting $\Phi[\cdot]$ denote the standard normal distribution function, at time t player i chooses strategy 1 with probability $\Phi[(3/2 - \alpha_t - \beta_t y_{t-1} - (1 - \beta_t)x_{it-1})/(\sigma_{\zeta_{it}} | x_{t-1})]$; strategy x , for $x = 2, \dots, 6$, with probability

$$\begin{aligned} & \Phi[(x + 1/2 - \alpha_t - \beta_t y_{t-1} - (1 - \beta_t)x_{it-1})/(\sigma_{\zeta_{it}} | x_{t-1})] \\ & - \Phi[(x - 1/2 - \alpha_t - \beta_t y_{t-1} - (1 - \beta_t)x_{it-1})/(\sigma_{\zeta_{it}} | x_{t-1})]; \end{aligned}$$

and strategy 7 with probability $1 - \Phi[(13/2 - \alpha_t - \beta_t y_{t-1} - (1 - \beta_t)x_{it-1})/(\sigma_{\zeta_{it}} | x_{t-1})]$.

The evolution of players' beliefs can be described by (5) and (6) as before, with players' strategy choices determined by rounding the x_{it} , y_t computed by evaluating $f(\cdot)$ at those rounded values, and the ζ_{it} satisfying the distributional assumptions in Section 3. The dynamics are still a Markov process with the state vector x_t representing players' beliefs; the only difference is in how players' beliefs determine their strategy choices and the summary statistic. To see how this difference affects the results, let $h(x_{1t}, \dots, x_{nt}) \equiv f(\hat{x}_{1t}, \dots, \hat{x}_{nt})$, where \hat{x}_{it} denotes the rounded value of x_{it} , so that $h(\cdot)$ determines the value of the summary statistic that players observe as a function of their beliefs. Because $h(a + bx_{1t}, \dots, a + bx_{nt}) = a + bh(x_{1t}, \dots, x_{nt})$ only by coincidence, Proposition 1 does not carry over directly to the present model. I now argue, however, that Propositions 2, 3, and 4—which do not depend on Proposition 1—continue to hold exactly as stated with discrete strategy choice, and that Propositions 5 and 6—which do depend on Proposition 1—hold approximately whenever the grid of feasible strategies is sufficiently fine.

The proof of Proposition 2 goes through because the x_{it} remain continuously variable and y_{t-1} , the only variable in (6) that is affected by rounding, cancels

²⁷ It would be easy to allow other values for the boundaries of these choice regions, or to estimate them. But rounding this way is optimal for the payoffs in most of VHBB's treatments when players are risk-neutral and x_{it} is viewed as player i 's estimate of $E_{t-1}(y_t | x_{t-1})$, and allowing other values is unlikely to change the results very much.

out of the Lyapunov function in (10). Thus, under the stated variance conditions and bounds on the β_t , players' beliefs and strategy choices still converge with probability one to one of the pure-strategy equilibria of the game. The proofs of Propositions 3 and 4 also go through, because $h(\cdot)$ inherits the weak monotonicity properties of $f(\cdot)$. Thus, the qualitative comparative dynamics are unaltered by discrete choice.

Discrete choice does affect the quantitative comparative dynamics. Proposition 1, and thus Propositions 5 and 6, remain formally valid, with the conditional distributions of the ζ_{it} adjusted for discreteness and μ_t , σ_t^2 , γ_t , and δ_{st} evaluated accordingly. The difficulty of evaluating those parameters limits the usefulness of these results. It is more helpful to identify settings in which the means and variances for the model with continuously variable strategies are good approximations to their analogs with discrete choice.

Plainly, the errors in approximating the means and variances in this way cannot be negligible unless the probability of significant boundary effects is negligible.²⁸ Given this, the errors are negligible whenever the grid of feasible strategies is sufficiently fine. To see this, note that rounding the x_{it} can alter their order only by creating ties, which does not affect order statistics. It follows that if $f(\cdot)$ is an order statistic, then $h(x_{1t}, \dots, x_{nt})$ equals the rounded value of $f(x_{1t}, \dots, x_{nt})$; and if $f(\cdot)$ is a convex combination of order statistics, then $h(x_{1t}, \dots, x_{nt})$ equals that combination of the rounded values of those order statistics, rounded as necessary when y_t is also required to be discrete. Thus, the effects of rounding in (5) and (6) are filtered through y_t , bounded by half the size of the grid each period, and additive across periods. It can then be shown, using the fact that the convergence of the x_{it} is asymptotically geometric in both models, that the cumulative effect of rounding on y_t remains finite with probability one and approaches zero with the size of the grid. The quality of the approximations reported in Crawford (1992) suggests that the grids in VHBB's experiments were fine enough to make the effects of the discreteness negligible.

This section's analysis shows that most of the theoretical problems caused by discrete strategy spaces can be overcome for the present model. It would, however, be significantly easier to analyze the results of experiments with approximately continuous strategy spaces.

6. ESTIMATION

Although the theoretical analysis of Sections 4 and 5 sheds considerable light on the patterns of discrimination among equilibria in VHBB's experiments, a full explanation depends on the values of the parameters left undetermined by the theory. The model provides a simple specification of the distribution of

²⁸ Boundary effects were a significant problem in both of VHBB's minimum treatments and, to a lesser extent, one of their median treatments, Ω .

players' beliefs within which those parameters can be estimated. This section reports estimates computed using VHBB's experimental data, in preparation for the analysis of coordination outcomes in Section 7.²⁹

The model was estimated separately for each of VHBB's experimental treatments that came first in its sequence, pooling the data from all such runs of each treatment.³⁰ The estimates were obtained by maximum-likelihood techniques, taking the discreteness of strategy choices into account using the ordered probit model of Section 5, with the latent variables determined as in (5) and (6) and the error terms taken to be heteroskedastic as in (4), conditionally independent across players, and conditionally normally distributed.³¹

The model was first estimated for each treatment with no restrictions on the behavioral parameters.³² Allowing the estimates to vary freely over time in this way maximizes the usefulness of their plausibility as an informal diagnostic and minimizes the risk of specification bias. However, as explained in Section 3, the large number of parameters (5 per period minus 3) also leaves room for doubt about whether the model can explain the dynamics only through ad hoc parameter variations. To resolve this doubt and to provide the structure needed for beyond-sample prediction, the model was reestimated for each treatment under the simplest intertemporal constraints that appeared to have a chance of not being rejected. These required that $\alpha_t = 0$ for $t = 1, \dots$ in the median treatments; that $\alpha_t = \alpha$ for $t = 1, \dots$ in the minimum treatments, in which $\alpha_t = 0$ is implausible; that $\beta_t = \beta$ for $t = 1, \dots$ in all treatments; and that the parameters of the variance-covariance matrix, $\sigma_{\varepsilon t}^2$, $\sigma_{\eta t}^2$, and κ_t , all decline over time like $1/t^\lambda$ for $t = 1, \dots$ in all treatments.³³ Including α_0 and $\sigma_{\varepsilon 0}^2$, which are allowed to vary independently to reflect the mean of players' initial beliefs and their different stochastic specification, these constraints reduce the number of inde-

²⁹ Broseta (1993b) estimates a related model that allows for autoregressive conditional heteroskedasticity in the idiosyncratic components of subjects' beliefs. This richer intertemporal stochastic structure describes the data better in some respects.

³⁰ An exception was made in including the random-pairing minimum treatment, which was always run following other treatments. Only the longer of the two runs in this treatment was used, to avoid problems with missing observations. All other treatments were sometimes run first in a series and sometimes later on; prior experience clearly affected subjects' initial beliefs, making pooling inadvisable. The two (out of nine) runs of the large-group minimum treatment in which entire strategy profiles were announced were omitted for similar reasons.

³¹ To make the computations manageable, the likelihood function was constructed period by period, using the \hat{x}_{it-1} , subjects' (rounded) strategy choices from the previous period, as proxies for the x_{it-1} , their unobservable (unrounded) beliefs. Without this substitution, this procedure would yield consistent and asymptotically efficient parameter estimates under my assumptions. The grid appears to have been sufficiently fine that the substitution made little difference to the estimates.

³² In each case the reported estimates are unrestricted maximum likelihood, except that β_1 in treatment Ω was constrained to lie between 0 and 1 because the unrestricted estimate, 1.28, was implausible. In this and a few other cases, indicated by dashes in parentheses in Tables I–V, it proved impossible to compute standard errors.

³³ No cross-treatment restrictions were imposed, because the unconstrained estimates suggest that most simple restrictions of this kind would be rejected, and such restrictions are neither suggested by the theory nor needed for the analysis.

TABLE I
ESTIMATES FOR MEDIAN TREATMENT I
(54 observations each period; asymptotic standard errors in parentheses)

t	α_t	β_t	Unconstrained Model					Log L
			$\sigma_{\epsilon t}^2$	$\sigma_{\eta t}^2$	κ_t	$\sigma_{z t}^2$	$\sigma_{\xi t}^2$	
0	4.75 (0.18)	—	1.62 (0.37)	—	—	1.62	1.62	-86.7
1	-0.07 (0.13)	0.56 (0.12)	0.57 (0.18)	0.19 (0.14)	0.21 (0.11)	1.17	0.86	-67.9
2	0.15 (0.11)	0.80 (0.14)	0.41 (0.11)	0.11 (0.10)	-0.12 (0.09)	0.58	0.53	-59.4
3	0.02 (0.07)	0.63 (0.12)	0.14 (0.04)	0.05 (0.08)	-0.01 (0.07)	0.25	0.17	-37.9
4	0.10 (0.10)	0.57 (0.20)	0.05 (0.02)	0.55 (0.27)	-0.16 (0.05)	0.22	0.18	-24.2
5	-0.10 (0.12)	0.48 (0.23)	0.04 (0.02)	0.50 (0.27)	0.14 (0.05)	0.20	0.14	-15.6
6	0.09 (0.09)	0.41 (—)	0.04 (0.02)	0.43 (0.02)	-0.04 (0.02)	0.12	0.05	-8.7

t	α_t	β_t	Constrained Model					$\Sigma_t \text{Log } L_t$
			$\sigma_{\epsilon t}^2$	$\sigma_{\eta t}^2$	κ_t	$\sigma_{z t}^2$	$\sigma_{\xi t}^2$	
0	4.75 (0.18)	—	1.62 (0.37)	—	—	1.62	1.62	-334.3
1	0.00	0.58 (0.07)	0.43 (0.20)	0.78 (0.31)	-0.05 (0.13)	0.82	0.54	
2	0.00	0.58 (0.07)	0.23 (0.06)	0.43 (0.11)	-0.03 (0.07)	0.61	0.47	
3	0.00	0.58 (0.07)	0.17 (0.03)	0.30 (0.08)	-0.02 (0.05)	0.44	0.33	
4	0.00	0.58 (0.07)	0.13 (0.02)	0.24 (0.07)	-0.01 (0.04)	0.30	0.22	
5	0.00	0.58 (0.07)	0.11 (0.02)	0.20 (0.07)	-0.01 (0.03)	0.21	0.16	
6	0.00	0.58 (0.07)	0.09 (0.02)	0.17 (0.07)	-0.01 (0.02)	0.16	0.12	

pendent behavioral parameters to 8 in the minimum treatments and 7 in the median treatments, independent of the number of periods.

Tables I–V report the unconstrained and constrained parameter estimates for the three median treatments, Γ , Ω , and Φ , and the large-group and random-pairing minimum treatments, A and C , with asymptotic standard errors in parentheses. Each treatment lasted 10 periods except for treatment C , in which the run for which the estimates were computed lasted 5 periods, but no estimates are reported for some periods near the ends of the median treatments in which there was no longer enough sample variation to identify the parameters. To avoid specification bias, the constrained estimates for treatment A were computed omitting the data from period 9, in which the unconstrained estimates revealed large end-of-treatment effects, described below. The estimates of the $\sigma_{z t}^2$ and $\sigma_{\xi t}^2$ reported for the median treatments were approximated

TABLE II
ESTIMATES FOR MEDIAN TREATMENT Ω
(27 observations each period; asymptotic standard errors in parentheses)

t	α_t	β_t	Unconstrained Model					$\text{Log } L$
			$\sigma_{\varepsilon t}^2$	$\sigma_{\eta t}^2$	κ_t	$\sigma_{z t}^2$	$\sigma_{\xi t}^2$	
0	6.26 (0.36)	—	2.37 (1.04)	—	—	2.37	2.37	−35.9
1	0.34 (0.30)	1.00 (—)	1.47 (0.87)	2.47 (1.73)	−1.90 (1.16)	7.00	7.00	−25.8
2	−0.16 (0.22)	0.95 (0.18)	0.04 (0.06)	0.15 (0.14)	0.04 (0.15)	1.08	1.06	−12.0
3	−0.01 (0.36)	0.97 (0.44)	0.02 (0.05)	0.03 (0.05)	−0.03 (0.88)	0.05	0.05	−0.2

t	α_t	β_t	Constrained Model					$\Sigma_t \text{Log } L_t$
			$\sigma_{\varepsilon t}^2$	$\sigma_{\eta t}^2$	κ_t	$\sigma_{z t}^2$	$\sigma_{\xi t}^2$	
0	6.26 (0.36)	—	2.37 (1.04)	—	—	2.37	2.37	−76.5
1	0.00	0.97 (0.13)	0.89 (0.41)	0.90 (0.64)	0.54 (0.60)	2.91	2.91	
2	0.00	0.97 (0.13)	0.09 (0.04)	0.09 (0.06)	0.06 (0.05)	0.35	0.35	
3	0.00	0.97 (0.13)	0.02 (0.16)	0.02 (0.02)	0.01 (0.01)	0.03	0.03	

TABLE III
ESTIMATES FOR MEDIAN TREATMENT Φ
(27 observations each period; asymptotic standard errors in parentheses)

t	α_t	β_t	Unconstrained Model					$\text{Log } L$
			$\sigma_{\varepsilon t}^2$	$\sigma_{\eta t}^2$	κ_t	$\sigma_{z t}^2$	$\sigma_{\xi t}^2$	
0	4.71 (0.20)	—	0.97 (0.30)	—	—	0.97	0.97	−38.5
1	−0.30 (0.12)	0.95 (0.16)	0.22 (0.10)	0.07 (0.11)	0.09 (0.10)	0.29	0.29	−23.6
2	0.11 (—)	0.40 (—)	0.00 (1.13)	0.00 (0.94)	0.00 (0.57)	0.11	0.00	−4.5
3	0.16 (0.18)	0.92 (0.27)	0.04 (0.04)	0.10 (—)	0.00 (—)	0.05	0.05	−5.1

t	α_t	β_t	Constrained Model					$\Sigma_t \text{Log } L_t$
			$\sigma_{\varepsilon t}^2$	$\sigma_{\eta t}^2$	κ_t	$\sigma_{z t}^2$	$\sigma_{\xi t}^2$	
0	4.71 (0.20)	—	0.97 (0.30)	—	—	0.97	0.97	−79.9
1	0.00	0.74 (0.11)	0.22 (0.10)	0.16 (0.14)	−0.01 (0.08)	0.43	0.36	
2	0.00	0.74 (0.11)	0.08 (0.02)	0.06 (0.06)	−0.00 (0.03)	0.13	0.10	
3	0.00	0.74 (0.11)	0.04 (0.02)	0.03 (0.03)	−0.00 (0.01)	0.05	0.04	

TABLE IV
ESTIMATES FOR LARGE-GROUP MINIMUM TREATMENT \mathcal{A}
(107 observations each period; asymptotic standard errors in parentheses)

t	α_t	β_t	Unconstrained Model		κ_t	Log L
			$\sigma_{\epsilon t}^2$	$\sigma_{\eta t}^2$		
0	5.45 (0.19)	—	3.47 (0.64)	—	—	—180.5
1	0.51 (0.23)	0.52 (0.09)	1.00 (0.41)	0.04 (0.03)	-0.21 (—)	—184.1
2	0.10 (0.41)	0.34 (0.14)	6.85 (2.69)	0.75 (0.32)	2.08 (0.93)	—170.4
3	0.27 (0.25)	0.54 (0.11)	2.96 (0.98)	0.48 (0.20)	1.01 (0.45)	—154.8
4	-0.63 (0.37)	0.15 (0.13)	6.00 (2.06)	0.52 (0.22)	1.64 (0.67)	—140.7
5	-0.54 (0.58)	0.20 (0.25)	1.76 (1.28)	0.76 (0.27)	1.06 (0.54)	—106.3
6	0.11 (0.22)	0.46 (0.22)	1.32 (0.51)	1.14 (0.45)	1.06 (0.40)	—115.5
7	-0.35 (0.35)	0.40 (0.18)	0.55 (0.42)	0.24 (0.11)	0.19 (0.19)	—78.3
8	-0.49 (0.25)	0.11 (0.17)	1.17 (0.40)	0.39 (0.19)	0.59 (0.24)	—80.7
9	-5.19 (1.97)	-4.47 (2.24)	26.0 (1.73)	37.7 (1.04)	30.5 (1.30)	—88.0

t	α_t	β_t	Constrained Model		κ_t	$\Sigma_t \text{Log } L_t$
			$\sigma_{\epsilon t}^2$	$\sigma_{\eta t}^2$		
0	5.45 (0.19)	—	3.47 (0.64)	—	—	—1265.0
1	-0.27 (0.09)	0.25 (0.04)	3.05 (0.61)	0.59 (0.14)	0.91 (0.25)	
2	-0.27 (0.09)	0.25 (0.04)	2.34 (0.35)	0.45 (0.09)	0.69 (0.17)	
3	-0.27 (0.09)	0.25 (0.04)	2.00 (0.27)	0.38 (0.08)	0.59 (0.14)	
4	-0.27 (0.09)	0.25 (0.04)	1.79 (0.23)	0.34 (0.07)	0.53 (0.12)	
5	-0.27 (0.09)	0.25 (0.04)	1.64 (0.21)	0.32 (0.06)	0.49 (0.11)	
6	-0.27 (0.09)	0.25 (0.04)	1.53 (0.21)	0.29 (0.06)	0.45 (0.10)	
7	-0.27 (0.09)	0.25 (0.04)	1.44 (0.20)	0.28 (0.06)	0.43 (0.10)	
8	-0.27 (0.09)	0.25 (0.04)	1.37 (0.20)	0.26 (0.06)	0.41 (0.09)	

from the estimated parameters using (23)–(25) and assuming normality of the z_{it} ; large boundary effects made it impractical to compute the analogous estimates for the minimum treatments.

Despite its simplicity, the model gives a reasonably accurate statistical summary of subjects' behavior in all three median treatments. The unconstrained parameter estimates are generally plausible and reasonably stable across periods. With minor exceptions the trend parameters, α_t for $t = 1, \dots$, are not

TABLE V
ESTIMATES FOR RANDOM-PAIRING MINIMUM TREATMENT C
(16 observations each period; asymptotic standard errors in parentheses)

t	α_t	β_t	Unconstrained Model		κ_t	Log L
			$\sigma_{\epsilon t}^2$	$\sigma_{\eta t}^2$		
0	4.30 (1.13)	—	17.29 (10.47)	—	—	−28.2
1	1.59 (0.71)	0.67 (0.28)	5.26 (3.15)	0.96 (0.49)	2.24 (1.15)	−22.6
2	0.35 (0.61)	0.51 (0.65)	4.74 (2.24)	3.33 (2.87)	3.97 (2.30)	−24.2
3	1.50 (0.01)	1.00 (0.01)	9.30 (6.48)	9.30 (6.48)	9.30 (6.48)	−17.0
4	0.80 (0.91)	0.42 (0.37)	6.51 (4.16)	0.72 (0.44)	2.16 (1.32)	−18.8

t	α_t	β_t	Constrained Model		κ_t	$\Sigma_t \text{Log } L_t$
			$\sigma_{\epsilon t}^2$	$\sigma_{\eta t}^2$		
0	4.30 (1.13)	—	17.29 (11.02)	—	—	−119.5
1	1.21 (0.44)	0.54 (0.17)	8.30 (4.89)	1.01 (0.57)	2.86 (1.68)	
2	1.21 (0.44)	0.54 (0.17)	8.04 (2.56)	0.98 (0.35)	2.77 (0.94)	
3	1.21 (0.44)	0.54 (0.17)	7.89 (2.64)	0.96 (0.39)	2.72 (1.00)	
4	1.21 (0.44)	0.54 (0.17)	7.79 (3.37)	0.95 (0.49)	2.68 (1.26)	

significantly different from 0. The adjustment parameters, β_t , are significantly different from 0, ranging from around 0.5 in treatment Γ to just below 1.0 in treatment Ω . The initial mean, α_0 , is noticeably higher in treatment Ω , in which (unlike in treatment Γ) the cost of failing to coordinate on the highest effort is no greater than for lower efforts, and the simplicity of the payoff function may make the argument for the highest effort easier to apprehend.³⁴ However, α_0 is not significantly higher in treatment Γ , where higher efforts are also associated with greater efficiency, than in treatment Φ , where they are not; and the other α_t are no higher in treatments Γ and Ω than in treatment Φ . In treatment Ω the α_t and β_t are close to the values that would correspond to best-reply dynamics if there were no dispersion, but in the other treatments the α_t and β_t are significantly different from both fictitious play and best-reply dynamics.

With minor exceptions, the unconstrained estimates of the variances $\sigma_{\epsilon t}^2$ and $\sigma_{\eta t}^2$ decline smoothly toward 0 over time, except for an upward jump in $\sigma_{\eta t}^2$ from

³⁴ Median treatment Γ combined a common preference for a higher median, other things equal, with increasingly severe penalties for being further and further away from the median. Median treatment Φ maintained these penalties while eliminating the preference for a higher median; and median treatment Ω maintained the preference for a higher median while imposing different penalties for being away from the median, usually higher than in treatment Γ but independent of the distance. See VHBB (1991) or Crawford (1991, Section 2) for details.

period 3 to period 4 in treatment Γ . The estimated covariances, κ_t , are almost all insignificantly different from 0. The implied values of the $\sigma_{\xi t}^2$ and $\sigma_{\eta t}^2$ also decline smoothly toward 0, except for a sharp upward jump from period 0 to period 1 in treatment Ω . Inspecting the data for the periods beyond those for which estimates are reported reveals that maximum likelihood estimates of $\sigma_{\xi t}^2$ and $\sigma_{\eta t}^2$ would differ at most slightly from 0, continuing their generally downward trends and implying a continuing downward trend in the $\sigma_{\xi t}^2$ and $\sigma_{\eta t}^2$.

The standard errors of both the unconstrained and the constrained estimates indicate that constraints that rule out strategic uncertainty, $\sigma_{\xi t}^2 \equiv \sigma_{\eta t}^2 \equiv \kappa_t \equiv 0$, would be strongly rejected in each case. The impression that the dispersion of subjects' beliefs diminished steadily over time created by the unconstrained estimates is confirmed by the estimates of λ , the common rate of decline of $\sigma_{\xi t}^2$, $\sigma_{\eta t}^2$, and κ_t in the constrained specification, which are significantly positive in all three median treatments: 0.86 (0.34) in Γ , 3.27 (0.77) in Ω , and 1.52 (0.53) in Φ (standard errors in parentheses). These estimates do not fully explain the strong convergence VHBB observed in the median treatments, because in treatments Γ and Φ they are not significantly greater than one as Proposition 2's variance condition requires in this case. This might reflect the fact that Proposition 2's condition is sufficient, but not necessary, or that convergence with probability one is an unrealistically stringent criterion.

Likelihood ratio tests show that the constraints cannot be rejected in treatments Ω and Φ , but are strongly rejected in treatment Γ . The χ^2 statistics are 5.3 with 10 degrees of freedom for treatment Ω , with p -value 0.87; 16.6 with 10 degrees of freedom for treatment Φ , with p -value 0.08; and 67.8 with 25 degrees of freedom for treatment Γ . Otherwise the constrained estimates are quite plausible, and similar to the unconstrained estimates.

The rejection in treatment Γ appears to be due to the upward jump in the dispersion of players' responses from period 3 to period 4, which is likely to violate any simple intertemporal restrictions. As the rejected constraints are neither an implication nor a necessary part of the theory, the rejection has no bearing on its validity. Because they can still be used to show that the model can explain the dynamics without ad hoc variations in the parameters, I have not tried to find alternative constraints that would not be rejected.

The model does less well summarizing subjects' behavior in the minimum treatments. The unconstrained model yields implausible estimates of the trend and adjustment parameters in period 9 of the large-group treatment, A , in which significant end-of-treatment effects were visible, with some subjects jumping all the way from effort 1 to effort 7.³⁵ The unconstrained estimates are

³⁵ These end effects seemed to be due to subjects' perception that the coordination of the timing of adjustments required to break out of an inefficient equilibrium was unlikely to be achieved in any period other than the last; see Crawford (1991, p. 57). They were strongest in the first three runs of treatment A , in which the endpoint was announced; but they were visible even in the last four runs, in which it was not announced, perhaps because subjects believed that the experimenters or their partners were habituated to the decimal system. Estimates computed separately for each of these regimes did not differ enough to justify reporting them separately.

otherwise plausible and reasonably stable over time. The initial mean, α_0 , is higher in treatment *A* than in the random-pairing treatment, *C*, and roughly comparable to the initial means estimated for the median treatments. The trend parameters, α_t for $t = 1, \dots$, are usually insignificantly different from 0 in treatment *A* and positive, often significantly so, in treatment *C*, showing some tendency to decline over time in treatment *A* but no clear trend in treatment *C*. The trend parameters are systematically higher in treatment *C* than in all other treatments, suggesting that subjects may have corrected for the fact that their current pair minimum tends to underestimate the median of the other subjects' efforts that determined their best replies (see footnote 10). The adjustment parameters, β_t , are usually significantly positive in treatment *A* and positive, often significantly so, in treatment *C*. In each case the α_t and β_t differ significantly from the values implied by fictitious play and best-reply dynamics.³⁶

In both minimum treatments the unconstrained model yields estimates of the variance parameters that show only a weak tendency to decline over time, with several upward jumps in the estimated values of the $\sigma_{\varepsilon t}^2$ and $\sigma_{\eta t}^2$. This impression is confirmed by the constrained estimates of λ , the rate of decline (standard errors in parentheses): 0.39 (0.11) in treatment *A*, significantly positive but smaller than in the median treatments; and 0.05 (0.59) in treatment *C*, insignificantly positive. The convergence observed in treatment *A* is explained in Section 7 by considerations other than those addressed by Proposition 2, so the fact that its estimated λ is significantly less than 1 is no cause for concern. The much smaller rate of decline in treatment *C* reflects its extremely slow convergence, which was to be expected, given the very noisy information its subjects received about the population median that determined their best replies (see footnote 11).

The unconstrained estimates of the κ_t are generally positive, usually significantly so, in treatment *A* and positive, but generally insignificantly so, in treatment *C*, as in the median treatments. The fact that the estimated covariances are significantly different from zero only in the large-group minimum treatment has a plausible explanation. Because all subjects are trying to forecast the same summary statistic based on the same experience, high values of ε_{it} tend to be associated with high values of η_{it} in treatment *A*, where the relevant summary statistic is the minimum, so that the typical subject's effort is above it, and $y_{t-1} - x_{it-1} < 0$ in (4). In the other treatments, by contrast, the relevant summary statistic is the median, so that as many subjects' efforts are above it as below it.

Likelihood ratio tests show that the constraints are strongly rejected in treatment *A* even when the last period is omitted, but that they cannot be rejected in treatment *C*. The relevant χ^2 statistics are 107.4 with 34 degrees of freedom for treatment *A*; and 17.4 with 14 degrees of freedom for treatment *C*,

³⁶ Boylan and El-Gamal (1993) analyzed grouped data from different random-pairing experiments, comparing randomly perturbed versions of fictitious play and best-reply dynamics. On the assumption that subjects' unperturbed responses all followed either one or the other of these rules, they found that the evidence favors fictitious play.

with p -value 0.24. The constrained estimates otherwise seem plausible, and are generally quite close to the unconstrained estimates.

Inspecting the unconstrained estimates suggests that the rejection in treatment A is again due to upward jumps in the dispersion of subjects' responses, which are likely to violate any simple intertemporal restrictions. As in treatment F , this has no bearing on the validity of the theory, and I shall use the rejected constraints as simplifying restrictions to show that the model can explain the dynamics without ad hoc parameter variations.

7. COORDINATION OUTCOMES

The parameter estimates can be used to infer the probability distributions of coordination outcomes from the results of a limited number of experimental trials. Those distributions can be approximated analytically as in Section 4, but they can be estimated precisely only by repeated simulation. This section describes simulations based on the constrained parameter estimates reported in Section 6 and uses the resulting estimated distributions to flesh out the explanation of VHBB's results outlined in Section 4.

The estimated distributions play an essential role in evaluating the model. As explained in Section 3, the model can explain the dynamics VHBB observed only if there are significant differences in players' beliefs. Because it does not seem useful to try to explain such differences when players are otherwise identical, the model treats them as error terms within a given stochastic structure, and therefore determines a probability distribution over outcomes rather than predicting a particular outcome. This inevitably limits the model's goodness of fit, so that it should be evaluated by comparing the frequency distributions from the experiments with the probability distributions it implies.

Such comparisons can be made either conditionally or unconditionally; both kinds of test are reported here. The simulations take into account the discreteness of the rounded y_t and x_{it} , denoted \hat{y}_t and \hat{x}_{it} . Because \hat{y}_t and the \hat{x}_{it} normally approach a common limit except in treatment C , I simplify by focusing on \hat{y}_t in treatments other than C .

I begin by describing the results of the conditional simulations and comparisons. These were designed to estimate the conditional probability distribution of \hat{y}_t in the last period for which estimates are reported, given the realized value of \hat{x}_{t-1} .³⁷ Because the realized value of \hat{x}_{t-1} differed across runs in each treatment, a separate estimate based on 500 simulation runs was computed for each. In order to test the model's ability to "predict" beyond sample, these simulations were based on parameter estimates computed under the intertemporal constraints discussed in Section 6 but omitting the last period of data in each treatment but A , and the last two periods of data in treatment A . (These estimates are not reported, but are close to the constrained estimates in Tables I–V.)

³⁷ To make the computations manageable, the \hat{x}_{it-1} were used as proxies for the unobservable x_{it-1} . This appears to have made little difference to the results.

In all twelve experiments with median games and all seven experiments with treatment *A*, the model assigns conditional probability of at least 0.998, given the realized value of \hat{x}_{t-1} , to the value of \hat{y}_t realized in the last period.³⁸ These fits are so close that formal tests are unnecessary. In the experiment with treatment *C* considered here the model assigns conditional probability 0.37 to 6, the realized last-period median (the relevant summary statistic here) and conditional probabilities of 0.53, 0.09, and 0.01 to the alternative values 7, 5, and 4.³⁹ The χ^2 statistic for a test of goodness of fit is 1.7 with 6 degrees of freedom, with *p*-value 0.94. (The analogous test for the \hat{x}_{it} is not independent of this test. However, its χ^2 statistic is 3.8 with 6 degrees of freedom, with *p*-value 0.70.)

I conclude that the model conditionally predicts the \hat{y}_t well in every treatment. These are weak tests, because it is not difficult to make conditional predictions of variables that vary as little over time as the \hat{y}_t do. But this weakness is inherent in the data set, which simply does not allow very powerful tests of this kind.⁴⁰

I now turn to the unconditional simulations, which were designed to estimate the prior probability distributions of the \hat{y}_t . These simulations were carried out using both the unconstrained and the constrained parameter estimates, with 500 runs for each treatment. I report here only the results for the constrained estimates, despite the fact that the constraints were rejected in treatments *I* and *A*, because they provide a more stringent test of the model's ability to explain the dynamics and help to answer the criticism that it can explain them only through ad hoc variation in its parameters. (The simulation results for the unconstrained estimates, which are very similar, are reported in Crawford (1992).)

Because the entire time paths of \hat{y}_t are important in discriminating among alternative explanations, the results of these simulations are reported for as many periods as the estimates permitted in each treatment. The first and second lines of the cells in Tables VI–Xa report the actual frequency distributions and the estimated prior probability distributions of the \hat{y}_t . The distribution for the last period should approximate the common limiting distribution of \hat{y}_t and the \hat{x}_{it} implied by the model in each treatment but *C*, in which convergence was slow. Table Xb supplements the results for the \hat{y}_t in treatment *C* reported in Table Xa with the analogous distributions for the \hat{x}_{it} .

³⁸ All 500 runs yielded the realized median or minimum in all but one *I* and one *Ω* experiment, in which all but one of 500 runs yielded the realized median.

³⁹ Ambiguities of the median due to the evenness of *n* in treatment *C* were resolved by splitting the weight of ambiguous observations between the two possible values.

⁴⁰ Breaking the sample earlier and trying to predict longer histories of the \hat{y}_t would yield little additional information because the \hat{y}_t do not change much even in the early periods, and would make it even more difficult to estimate the intertemporal relationships required for beyond-sample prediction. Comparing conditional predictions of the \hat{x}_{it} with the experimental results in the treatments other than *C* would amount to a joint test of the model's ability to conditionally predict the \hat{y}_t , which has already been tested, and the assumption maintained in the simulations that the (unrounded) x_{it} are conditionally normally distributed, which is a working hypothesis rather than a substantive implication of the model. (Inspecting the data nevertheless suggests that conditional normality is approximately satisfied.)

TABLE VI
DYNAMICS OF \hat{y}_t IN MEDIAN TREATMENT I'
(actual frequency distributions in first lines,
simulated frequency distributions in second lines)

	\hat{y}_0	\hat{y}_1	\hat{y}_2	\hat{y}_3	\hat{y}_4	\hat{y}_5	\hat{y}_6
7	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	0.00	0.00	0.00	0.00	0.00	0.00	0.00
6	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	0.06	0.09	0.10	0.11	0.11	0.11	0.11
5	0.50	0.50	0.50	0.50	0.50	0.50	0.50
	0.62	0.56	0.53	0.52	0.53	0.52	0.52
4	0.50	0.50	0.50	0.50	0.50	0.50	0.50
	0.31	0.31	0.33	0.32	0.32	0.32	0.32
3	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	0.01	0.03	0.03	0.04	0.04	0.04	0.04
2	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	0.00	0.00	0.01	0.01	0.01	0.01	0.01
1	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	0.00	0.00	0.00	0.00	0.00	0.00	0.00

The coarseness of experimental frequency distributions based on 1–7 trials makes them unlikely to resemble the probability distributions implied by the model closely. Formal tests suggest, however, that the model replicates the dynamics in every treatment. The most informative tests appear to be χ^2 tests of goodness of fit comparing the estimated prior distributions for the last-period \hat{y}_t with the corresponding experimental frequency distributions, because it is harder to track the prior distributions for the last period than for earlier

TABLE VII
DYNAMICS OF \hat{y}_t IN MEDIAN TREATMENT Ω
(actual frequency distributions in first lines,
simulated frequency distributions in second lines)

	\hat{y}_0	\hat{y}_1	\hat{y}_2	\hat{y}_3
7	0.67	0.67	0.67	0.67
	0.34	0.37	0.36	0.36
6	0.00	0.00	0.00	0.00
	0.55	0.47	0.47	0.47
5	0.33	0.33	0.33	0.33
	0.10	0.15	0.15	0.15
4	0.00	0.00	0.00	0.00
	0.01	0.02	0.02	0.02
3	0.00	0.00	0.00	0.00
	0.00	0.00	0.00	0.00
2	0.00	0.00	0.00	0.00
	0.00	0.00	0.00	0.00
1	0.00	0.00	0.00	0.00
	0.00	0.00	0.00	0.00

TABLE VIII
DYNAMICS OF \hat{y}_t IN MEDIAN TREATMENT Φ
(actual frequency distributions in first lines,
simulated frequency distributions in second lines)

	\hat{y}_0	\hat{y}_1	\hat{y}_2	\hat{y}_3
7	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00
6	0.00 0.04	0.00 0.05	0.00 0.05	0.00 0.05
5	0.67 0.68	0.67 0.66	0.67 0.66	0.67 0.66
4	0.33 0.28	0.33 0.29	0.33 0.29	0.33 0.29
3	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00
2	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00
1	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00

TABLE IX
DYNAMICS OF \hat{y}_t IN LARGE-GROUP MINIMUM TREATMENT A
(actual frequency distributions in first lines,
simulated frequency distributions in second lines)

	\hat{y}_0	\hat{y}_1	\hat{y}_2	\hat{y}_3	\hat{y}_4	\hat{y}_5
7	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00
6	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00
5	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00
4	0.29 0.09	0.00 0.01	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00
3	0.14 0.37	0.00 0.04	0.14 0.01	0.00 0.00	0.00 0.00	0.00 0.00
2	0.29 0.33	0.57 0.18	0.29 0.03	0.00 0.01	0.00 0.00	0.00 0.00
1	0.29 0.21	0.43 0.78	0.57 0.96	1.00 0.99	1.00 1.00	1.00 1.00

periods.⁴¹ These tests never reject in any treatment. The relevant χ^2 statistics, each with six degrees of freedom, are 2.54 in treatment C , with p -value 0.86; 0.00 in treatment A ; 1.55 in treatment Γ , with p -value 0.95; 2.95 in treatment Ω , with p -value 0.81; and 0.16 in treatment Φ , with p -value 0.99. (The

⁴¹ More powerful tests involve entire histories of \hat{y}_t and/or \hat{x}_{it} . These are complicated by statistical dependence, and the results suggest that rejection is unlikely.

TABLE Xa
DYNAMICS OF \hat{y}_t IN RANDOM-PAIRING MINIMUM TREATMENT C
(actual frequency distributions in first lines,
simulated frequency distributions in second lines)

	\hat{y}_0	\hat{y}_1	\hat{y}_2	\hat{y}_3	\hat{y}_4
7	0.00 0.06	0.00 0.09	0.00 0.18	1.00 0.30	0.00 0.38
6	0.00 0.14	0.00 0.21	0.00 0.26	0.00 0.27	1.00 0.28
5	0.00 0.25	1.00 0.29	1.00 0.28	0.00 0.26	0.00 0.22
4	1.00 0.29	0.00 0.27	0.00 0.21	0.00 0.14	0.00 0.10
3	0.00 0.19	0.00 0.11	0.00 0.06	0.00 0.04	0.00 0.01
2	0.00 0.08	0.00 0.02	0.00 0.01	0.00 0.00	0.00 0.00
1	0.00 0.01	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00

analogous tests for earlier periods and/or the unconstrained estimates, which are not independent, would yield the same conclusions, with the exception of the constrained estimates for period 2 in treatment A.) The model replicates the unconditional frequency distributions from VHBB's experiments.

The statistical analysis suggests that VHBB's results were not anomalous. In some cases, however, the model strengthens or modifies the impressions created by the raw data. The inefficient outcomes they observed in treatment A now appear inevitable, appearing in all 500 simulation runs. It appears, however,

TABLE Xb
DYNAMICS OF THE \hat{x}_{it} IN RANDOM-PAIRING MINIMUM TREATMENT C
(actual frequency distributions in first lines,
simulated frequency distributions in second lines)

	\hat{x}_{i0}	\hat{x}_{i1}	\hat{x}_{i2}	\hat{x}_{i3}	\hat{x}_{i4}
7	0.31 0.30	0.31 0.33	0.25 0.37	0.63 0.42	0.50 0.46
6	0.00 0.09	0.06 0.09	0.19 0.10	0.00 0.11	0.00 0.11
5	0.13 0.09	0.31 0.11	0.19 0.11	0.19 0.10	0.25 0.10
4	0.19 0.09	0.06 0.11	0.06 0.11	0.06 0.10	0.06 0.09
3	0.06 0.09	0.06 0.10	0.06 0.10	0.00 0.08	0.00 0.07
2	0.06 0.08	0.06 0.07	0.13 0.06	0.13 0.06	0.13 0.05
1	0.25 0.25	0.13 0.20	0.13 0.16	0.00 0.13	0.06 0.12

that players might do better than VHBB's subjects in treatment Γ , and worse in treatment Ω . The model also suggests somewhat more strongly than the data that treatment C is likely to yield a high-effort equilibrium.

The model attributes the complete lack of history-dependence in treatment A to the strongly negative drift in this treatment, which overwhelmed its variance so that the lower bound on players' strategies determined the final outcome with high probability. The strong history-dependence in the median treatments and the moderate history-dependence in treatment C can be traced to the low to moderate drifts in these treatments, the large conditional variances of y_i in treatment C , and the small conditional variances of y_i in the median treatments. The analysis suggests, however, that the perfect history-dependence VHBB observed in all twelve of their experiments with median games overstates the importance of historical accidents in such environments. In the simulation results the correlation between \hat{y}_0 and the last \hat{y}_i reported is only 0.68 in treatment Γ , 0.76 in treatment Ω , and 0.90 in treatment Φ , so that the initial median "explains" only 46%, 58%, and 81% of the variances of the limiting outcomes in these treatments.

8. CONCLUSION

The model provides a simple, unified explanation for the complex patterns of history-dependence and discrimination among equilibria in VHBB's experiments. Three directions for further research seem promising.

It would be desirable to conduct similar experiments, in order to test the theory's predictions beyond the samples that influenced its specification, and to learn more about how the strategic environment influences the dispersion of players' beliefs and the rate at which it is eliminated as they accumulate forecasting experience.

It would also be of interest to relax the assumption that players' beliefs and strategy choices are statistically identical *ex ante*. The substance of this assumption is the absence of externally observable differences between players. It may be tractable to allow, instead, a fixed population made up of two or more observable "types" of players, with each player's payoffs and best replies determined in each play by his type, his strategy choice, and the current population distribution of strategy choices by type. A random-pairing model in which players' types are identified with their roles in a "divide-the-dollar" game, as in Young (1993b), seems of particular interest. Such a model, in which players' preferences over equilibria are opposed, is a natural complement to the present analysis and may help in interpreting experimental evidence on bargaining outcomes (see Roth (1987)).

Finally, it would be of interest to extend the analysis from *tacit* coordination, in which players communicate only by playing the game, to *explicit* coordination, in which players can send signals that are not directly related to the game (and which may not be costly). Tacit coordination raises many issues that must be resolved to understand explicit coordination, and generalizing the techniques developed here may yield useful models of explicit coordination. A natural

place to begin is games with one or more rounds of costless, simultaneous pre-play communication (see Farrell (1987), Palfrey and Rosenthal (1991), and Crawford (1990)) or the games with costly pre-play communication used in VHBB's (1993) experiments (see Crawford and Broseta (1994)). Considering how the meanings of players' messages and their decisions evolve may resolve the ambiguity that plagues traditional analyses of preplay communication and shed new light on proposed equilibrium refinements.

Dept. of Economics, University of California, San Diego, 9500 Gilman Dr., La Jolla, CA 92093-0508, U.S.A.

Manuscript received January, 1992; final revision received May, 1994.

REFERENCES

- BANKS, JEFFREY, CHARLES PLOTT, AND DAVID PORTER (1988): "An Experimental Analysis of Unanimity in Public Goods Provision Mechanisms," *Review of Economic Studies*, 55, 301–322.
- BOYLAN, RICHARD, AND MAHMOUD EL-GAMAL (1993): "Fictitious Play: A Statistical Study of Multiple Economic Experiments," *Games and Economic Behavior*, 5, 205–222.
- BROSETA, BRUNO (1993a): "Strategic Uncertainty and Learning in Coordination Games," Discussion Paper 93-34, University of California, San Diego.
- (1993b): "Estimation of a Game-Theoretic Model of Learning: An Autoregressive Conditional Heteroskedasticity Approach," Discussion Paper 93-35, University of California, San Diego.
- BRYANT, JOHN (1983): "A Simple Rational Expectations Keynes-Type Model," *Quarterly Journal of Economics*, 98, 525–528.
- COOPER, RUSSELL, DOUGLAS DEJONG, ROBERT FORSYTHE, AND THOMAS ROSS (1990): "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, 80, 218–233.
- COOPER, RUSSELL, AND ANDREW JOHN (1988): "Coordinating Coordination Failures in Keynesian Models," *Quarterly Journal of Economics*, 103, 441–463.
- CRAWFORD, VINCENT P. (1989): "Learning and Mixed-Strategy Equilibria in Evolutionary Games," *Journal of Theoretical Biology*, 140, 537–550; to be reprinted in *Proceedings of the Second Workshop on Knowledge, Belief, and Strategic Interaction*, ed. by Cristina Bicchieri and Brian Skyrms. New York: Cambridge University Press.
- (1990): "Explicit Communication and Bargaining Outcomes," *American Economic Review Papers and Proceedings*, 80, 213–219.
- (1991): "An 'Evolutionary' Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games and Economic Behavior*, 3, 25–59.
- (1992): "Adaptive Dynamics in Coordination Games," Discussion Paper 92-02R, University of California, San Diego.
- CRAWFORD, VINCENT P., AND BRUNO BROSETA (1994): "What Price Coordination? Auctioning the Right to Play as a Form of Preplay Communication," manuscript in preparation, University of California, San Diego.
- CRAWFORD, VINCENT P., AND HANS HALLER (1990): "Learning How to Cooperate: Optimal Play in Repeated Coordination Games," *Econometrica*, 58, 571–595.
- FARRELL, JOSEPH (1987): "Cheap Talk, Coordination, and Entry," *Rand Journal of Economics*, 18, 34–39.
- FOSTER, DEAN, AND H. PEYTON YOUNG (1990): "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology*, 38, 219–232.
- FREIDLIN, MARK, AND ALEXANDER WENTZELL (1984): *Random Perturbations of Dynamical Systems*. New York: Springer-Verlag.
- HARSANYI, JOHN, AND REINHARD SELTEN (1988): *A General Theory of Equilibrium Selection in Games*. Cambridge, Massachusetts: MIT Press.

- HARVEY, ANDREW (1989): *Forecasting, Structural Times Series Models and the Kalman Filter*. Cambridge, U.K., and New York: Cambridge University Press.
- ISAAC, R. MARK, D. SCHMIDTZ, AND J. WALKER (1989): "The Assurance Problem in a Laboratory Market," *Public Choice*, 62, 217–236.
- JONES, HOWARD (1948): "Exact Lower Moments of Order Statistics in Small Samples from a Normal Distribution," *Annals of Mathematical Statistics*, 19, 270–273.
- KANDORI, MICHIOH, GEORGE MAILATH, AND RAFAEL ROB (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61, 29–56.
- KIM YONG-GWAN (1990): "Evolutionary Analysis of Two-Person Finitely Repeated Coordination Games," manuscript, University of California, San Diego.
- LJUNG, LENNART, AND TORSTEN SODERSTROM (1983): *Theory and Practice of Recursive Identification*. Cambridge, Massachusetts: M.I.T. Press.
- MAYNARD SMITH, JOHN (1982): *Evolution and the Theory of Games*. Cambridge, U.K.: Cambridge University Press.
- McFADDEN, DANIEL (1984): "Econometric Analysis of Qualitative Response Models," Ch. 24 in *Handbook of Econometrics, Volume II*, ed. by Zvi Griliches and Michael Intriligator. Amsterdam and New York: Elsevier Science Publishing.
- McKELVEY, RICHARD, AND THOMAS PALFREY (1992): "An Experimental Study of the Centipede Game," *Econometrica*, 60, 803–836.
- NEVEL'SON, M. B., AND R. Z. HAS'MINSKII (1973): *Stochastic Approximation and Recursive Estimation*, Vol. 47, Translations of Mathematical Monographs. Providence, Rhode Island: American Mathematical Society.
- PALFREY, THOMAS, AND HOWARD ROSENTHAL (1991): "Testing for Effects of Cheap Talk in a Public Goods Game with Private Information," *Games and Economic Behavior*, 3, 183–220.
- ROBLES, JACK (1994): "Evolution and Long Run Equilibria in Coordination Games with Summary Statistic Payoff Technologies," Discussion Paper 94-12, University of California, San Diego.
- ROTH, ALVIN (1987): "Bargaining Phenomena and Bargaining Theory," Chapter 2 of *Laboratory Experimentation in Economics*, ed. by Alvin Roth. New York: Cambridge University Press.
- ROTH, ALVIN, AND IDO EREV (1995): "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior*, 8, forthcoming.
- ROTH, ALVIN, AND FRANCOISE SCHOUMAKER (1983): "Expectations and Reputations in Bargaining: An Experimental Study," *American Economic Review*, 73, 362–372.
- ROUSSEAU, JEAN-JACQUES (1973): "A Discourse on the Origin of Inequality," in *The Social Contract and Discourses* (translated by G. D. H. Cole). London: J. M. Dent & Sons, Ltd., 27–113.
- SCHELLING, THOMAS (1960): *The Strategy of Conflict*, First Edition. Cambridge, Massachusetts: Harvard University Press.
- TEICHROEW, D. (1956): "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*, 27, 410–426.
- VAN HUYCK, JOHN, RAY BATTALIO, AND RICHARD BEIL (1990): "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80, 234–248.
- (1991): "Strategic Uncertainty, Equilibrium Selection Principles, and Coordination Failure in Average Opinion Games," *Quarterly Journal of Economics*, 106, 885–910.
- (1993): "Asset Markets as an Equilibrium Selection Mechanism: Coordination Failure, Game Form Auctions, and Tacit Communication," *Games and Economic Behavior*, 5, 485–504.
- WOODFORD, MICHAEL (1990): "Learning to Believe in Sunspots," *Econometrica*, 58, 277–307.
- YOUNG, H. PEYTON (1993a): "The Evolution of Conventions," *Econometrica*, 61, 57–84.
- (1993b): "An Evolutionary Model of Bargaining," *Journal of Economic Theory*, 59, 145–168.