# TRADITIONAL, EVOLUTIONARY, AND ADAPTIVE APPROACHES TO EQUILIBRIUM SELECTION

**Vincent Crawford**
**University of California, San Diego**
**DELTA, Paris, May 1998**

## Introduction

● The lectures will discuss the leading alternative theoretical approaches to analyzing strategic behavior and equilibrium selection, focusing on games with complete (i.e. symmetric) information. There are two main goals: (i) to introduce the leading approaches and the modeling issues they address; and (ii) to examine their performance in the light of experimental evidence, in the hope of moving closer to the kind of understanding needed for applications. Six approaches will be considered (some only in passing):

1. Traditional equilibrium analysis and refinements, including Harsanyi and Selten's general theory and its notions of risk- and payoff-dominance

2. Equilibrium analyses of perturbed games

3. Rational learning models

4. Deterministic evolutionary dynamics

5. Long-run equilibria of stochastic evolutionary dynamics

6. Adaptive learning models

● The underlying assumptions of these approaches differ in three main ways:

a. *strategic sophistication*--the extent to which players' beliefs and behavior reflect an analysis of their environment as a game, taking its structure and other players' incentives into account

b. *strategic uncertainty*--the extent of players' uncertainty about others' strategies

c. the extent to which they seek to predict behavior entirely by theory, without using empirical information

● Ultimately, I believe that these differences can only be resolved by combining theory with empirical evidence. Here I focus on a particularly compelling body of evidence from experiments conducted by John Van Huyck, Ray Battalio, and Richard Beil ("VHBB"); and for concreteness I organize the discussion of theory as much as possible around their games:

VHBB, "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *AER* 1990

VHBB, "Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games," *QJE* 1991

VHBB, "Asset Markets as an Equilibrium Selection Mechanism: Coordination Failure, Game Form Auctions, and Tacit Communication," *GEB* 1993

● Much of my discussion follows

Crawford, "Learning Dynamics, Lock-in, and Equilibrium Selection in Experimental Coordination Games," to appear in Pagano and Nicita (eds.), *The Evolution of Economic Diversity*

which builds on the analyses in

Crawford, "Adaptive Dynamics in Coordination Games," *Econometrica* 1995

Bruno Broseta, "Strategic Uncertainty and Learning in Coordination Games" and "Estimation of a Game-Theoretic Model of Learning: An Autoregressive Conditional Heteroskedasticity Approach," UCSD Discussion Papers 93-34 and 93-35, 1993

Crawford and Broseta, "What Price Coordination? The Efficiency-Enhancing Effect of Auctioning the Right to Play," *AER* March 1998.

● A more general discussion of the use of experiments to learn about game theory can be found in

Crawford, "Theory and Experiment in the Analysis of Strategic Interaction," 206-242 in Kreps and Wallis (eds.), *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress*, Vol. I, Cambridge University Press, 1997

## VHBB's 1990 and 1991 experimental designs

● The designs have repeated play (usually for 10 periods) of symmetric coordination "stage games" in populations of subjects, interacting all at once as in evolutionary games against the field ("large groups") or in pairs drawn randomly from the population ("random pairing").

● Subjects chose simultaneously among 7 efforts, with payoffs and ex post optimal choices determined by own efforts and an order statistic, the population median or minimum effort in large groups or the current pair's minimum with random pairing.

● There were five leading treatments, varying the order statistic (minimum in 1990, median in 1991), the size of the subject population, and the patterns in which they interact (in 1990 minimum games were played either by the entire population of 14-16 or by random pairs, in 1991 median games were played by the entire population of 9); each population was large enough to make subjects treat their own influences on the order statistic as negligible.

● Explicit communication was prohibited throughout, the order statistic was publicly announced after each play (with random pairs told only pair minima), and the structure was publicly announced at the start, so subjects were uncertain only about others' efforts.

● The stage games all have seven strict, symmetric, Pareto-ranked equilibria, with players' best responses an order statistic of population efforts (even in the random-pairing minimum treatment, where for VHBB's payoffs, it's best, ex ante, to match the population median).

● The games have strategic structures somewhat like a faculty meeting that can't start until a given quorum is achieved--100% in the large-group minimum game, 50% in the large-group median games.

● The large-group minimum game, in particular, is like a more complex version of the Stag Hunt parable of the choice between autarky and a highly organized society whose productivity depends on coordination (a stylized version of a passage in Rousseau's *Discourse on Inequality*):

a. two efforts, 1 and 2, and two symmetric pure-strategy equilibria, "all-1" and "all-2"

b. all-2 is best for all, but the tension between its higher payoff and the risk of lower payoff if not all play 2 makes it less robust to strategic uncertainty

c. as a result, most people play 2 in small groups but 1 in large groups

**Minimum Effort**

| | | **2** | **1** |
|---|---|---|---|
| **Player'** | **2** | **2** | **0** |
| **Effort** | **1** | **1** | **1** |

**Stag Hunt**

**Advantages of VHBB's 1990 and 1991 designs**

● VHBB's designs pose the problem of coordination in a simple, tractable form, with a range of equilibria and a natural measure of their efficiency.

● They are canonical models of the emergence of conventions to solve coordination problems, which provide important evidence on how the difficulty of coordination depends on the number of players and the robustness of desirable outcomes.

● Their structures resemble economic important models, including Keynes' analogy between the stock market and newspaper beauty contests; macroeconomic "coordination failure" models; and team production models (assembly line, faculty meeting).

● The structural similarities across VHBB's treatments allow sharp tests of traditional refinements; and their large strategy spaces and variety of interaction patterns yield rich dynamics, which discriminate sharply between different models of learning and "evolution".

**VHBB's 1990 and 1991 results**

● The five leading treatments all evoked similar initial responses.

● Subjects almost always converged to an equilibrium of the stage game.

● But the dynamics varied with the treatment variables (order statistic, number of players, interaction pattern), with large differences in drift, history-dependence, rate of convergence, equilibrium selection:

a. in 12 out of 12 large-group median trials, there was near-perfect "lock-in" on the initial median (even though it varied across runs and was usually inefficient)

b. in 9 out of 9 large-group minimum trials, there was very strong downward drift, with subjects always approaching the least efficient equilibrium

c. in 2 out of 2 random-pairing minimum trials, there was very slow convergence, no discernible drift, and moderate inefficiency

● Comparing (a) and (b) reveals an "order statistic" or "robustness" effect, with coordination less efficient the smaller the groups that can disrupt desirable outcomes (Schelling's "three dimes" example).

● Comparing (b) and (c) reveals a "group size" effect, in which coordination is less efficient in larger groups (holding the order statistic constant, measured from the "bottom").

**VHBB's 1993 design**

● The design was the same as their 1990 and 1991 designs, with one treatment: repeated play (10-15 periods) of stage game consisting of one of the 1991 median games with the right to play it auctioned each period in a population of 18 (English clock auction, same price paid by all 9 winners).

● The market-clearing price was publicly announced after each period's auction, the median was publicly announced after each period's play, and the structure was publicly announced at the start.

● The stage game has a range of symmetric equilibria, in which all players bid the payoff of some equilibrium of the median game and then play that equilibrium, unless others bid differently; these equilibria are all subgame-perfect, all consistent with forward induction (iterated dominance).

**Advantages**

● The 1993 design is a "general equilibrium" analog of 1991 design, in which subjects choose among median games with the market-clearing price like an opportunity cost (regional or sectoral coordination failure models).

● The auction is an interesting form of preplay communication, in which subjects' willingnesses to pay may signal how they expect to play and thereby alleviate the tension due to strategic uncertainty.

● The design is a tractable example of learning to play extensive-form game.

**VHBB's 1993 results**

● In 8 of 8 trials, subjects quickly bid the price to a level that could only be recouped in the most efficient equilibrium and then converged to that equilibrium; the results give strong, precise selection among a wide range of subgame-perfect, forward induction equilibria.

● Auctioning the right to play has a strong efficiency-enhancing effect, a new and possibly important mechanism by which competition promotes efficiency; this effect is even stronger than suggested by the sum of the "order statistic," "optimistic subjects," and "forward induction" intuitions.

**A Portmanteau model**

● The model nests leading approaches to equilibrium selection in generalizations of VHBB's 1990, 1991 environments, so that analysis and econometric estimates can distinguish among them.

● In its most general form the model is an adaptive learning model with heterogeneous beliefs; this reflects the conviction that differences in players' beliefs cannot (usefully) be explained by theory alone because players have identical roles and preferences, and almost identical information.

● The model assumes that players ignore their individual influences on the order statistic, $y_t$, learn to predict it, and choose their optimal efforts, $x_{it}$.

● Learning is characterized in the style of the adaptive control literature, with players' beliefs represented by their optimal efforts, $x_{it}$, which are assumed to be continuously variable (this is relaxed in the papers, in which the $x_{it}$ are latent variables in an ordered probit model of discrete effort choice).

● The form of the learning rules and the "evolutionary" structure of VHBB's designs allow a simple statistical characterization of the dynamics of players' beliefs and efforts.

● $y_t$ can be written as a continuous function of the $x_{it}$:

(1) $$y_t \equiv f(x_{it},\ldots,x_{nt}),$$

where for any $x_{1t},\ldots x_{nt}$ and constants $a$ and $b \geq 0,$

(2) $$f(a + bx_{1t},\ldots, a + bx_{nt}) \equiv a + bf(x_{1t},\ldots, x_{nt}).$$

● The initial $x_{it}$ are i.i.d. draws from same population, with mean $\alpha_0$ and shocks $\zeta_{i0}$:

(3) $$x_{i0} = \alpha_0 + \zeta_{i0}.$$

● From then on the $x_{it}$ adjust toward the value suggested by the latest $y_{t\text{-}1}$:

(4) $$x_{it} = \alpha_t + \beta_t y_{t-1} + (1-\beta_t)x_{it-1} + \zeta_{it}, \quad t = 1,\ldots$$

● (4) suggests partial adjustment, but is best thought of as full adjustment to players' estimates of their optimal efforts, which respond less than fully to new $y_t$ observations because they are only part of players' information; (4) generalizes fictitious play and best-response adjustments, providing a robust approach to players' estimation problems.

● $\alpha_t$ and $\beta_t$ are exogenous behavioral parameters, with $0 < \beta_t \leq 1$ and $\alpha_t \to 0$ as $t \to \infty$.

● $\zeta_{it}$ are i.i.d. shocks, which represent differences in players' initial beliefs and interpretations of new observations; zero means by construction; variances $\sigma^2_{\zeta t}$ represent the level of strategic uncertainty; note how the "evolutionary" structure allows a simple statistical characterization of strategic uncertainty, with players' beliefs as i.i.d. draws from common population.

**Salient differences from other models**

● Equilibrium is allowed but not imposed, even in perturbed versions of the game; players are rational, but with heterogeneous beliefs and efforts.

● The model focuses on players' uncertainty about others' efforts and how they adjust, rather than uncertainty about others' preferences or their ability to implement decisions.

● Players' learning rules have the best-response structure built in, unlike "reinforcement learning"; subjects seemed to understand the structure, and such rules seem to work better here (compare Roth and Erev).

● Players' adjustments have time-varying, idiosyncratic degrees of inertia.

● The model's transition probabilities are generally positive but may decline over time, allowing nonergodic dynamics and eventual lock-in on a particular stage-game equilibrium.

**Preliminary analysis**

● The model is a Markov process with state vector $x_t$ and nonstationary transition probabilities, whose long-run steady states coincide with pure-strategy stage-game equilibria.

● Its recursive structure and i.i.d. shocks rule out unmodeled coordination (as by deduction); coordination can occur only via independent responses to common observations of $y_t$.

● The leading approaches are distinguished by different parameter values.

1. Equilibrium analyses have $\zeta_{it} \equiv \sigma^2_{\zeta t} \equiv 0$ for all $i$ and $t$ and $\alpha_t \equiv 0$ for all $t = 1, \ldots$ (because $x_{it} \equiv y_t$ when $\zeta_{it} \equiv \sigma^2_{\zeta t} \equiv 0$, $\beta_t$ is irrelevant here)

● Even with refinements, most such analyses do not discriminate among the multiple strict equilibria in VHBB's 1990, 1991 games or the multiple subgame-perfect, forward induction equilibria in VHBB's 1993 game; Harsanyi and Selten's general theory of equilibrium selection does discriminate, but makes systematic errors in both mean and dispersion.

2. Equilibrium analyses of perturbed games do not fit into the portmanteau

● Such analyses have problems like equilibrium analyses, but more subtle.

3. Rational learning models have $\zeta_{it} \equiv \sigma^2_{\zeta t} \equiv 0$ for all $i$ and $t$ and allow time-varying $\alpha_t$

● Such analyses do no better than equilibrium analyses ($\pi$ example).

4. Deterministic evolutionary dynamics have $\sigma^2_{\zeta 0} > 0$ , $\zeta_{it} \equiv \sigma^2_{\zeta t} \equiv \alpha_t \equiv 0$ for all *i* and *t* = 1, . . . ; thus they allow unlimited initial heterogeneity but no subsequent differences in players' interpretations of new observations

● Such dynamics imply history-dependent equilibrium selection and capture some effects of strategic uncertainty (the order statistic and players' interaction patterns affect the sizes of basins of attraction); but the initial state is taken as given, and the dynamics rule out "tunneling" across basins.

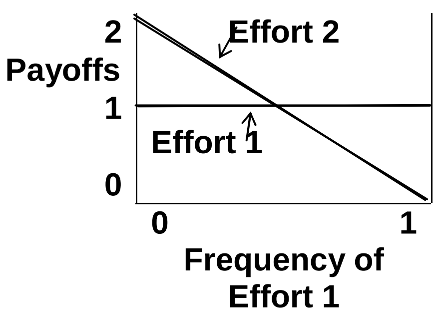


**Random pairing Stag Hunt**

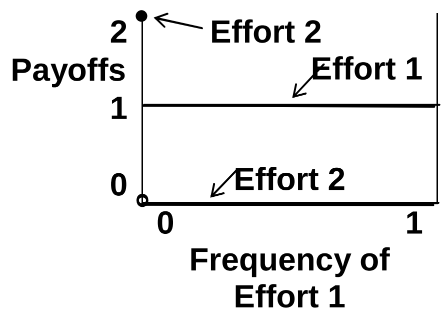


**Stag Hunt against the field**

● Proposition 1 gives a more general account of this history-dependence, showing how to use the model as an accounting system for keeping track of probabilities of changes in $y_t$**:**

**Proposition 1**: Suppose $\alpha_t = 0$ and $\beta_t \in (0,1]$ for all *t* = 1, . . . , and that there exists an integer *T* ≥ 1 such that $\zeta_{it} \equiv \sigma^2_{\zeta t} \equiv 0$ for all *t* = *T*, . . . . Then for all *i*, $x_{it} \to y_{T-1}$ monotonically, without overshooting, and $y_t = y_T$ for all *t* = *T* , . . . , independent of the number of players *n* and the order statistic $f(\cdot)$.

● The proof uses the fact that $y_t$ can change only if more players overshoot in one direction than in the other. (4) with $\alpha_t = 0$, $\beta_t \in (0,1]$ for *t* = 1, . . . , and $\zeta_{it}$ = 0 for *t* = *T*, . . . , implies that $x_{iT} - y_{T-1}$ always has the same sign as $x_{iT-1} - y_{T-1}$, with $x_{iT}$ closer to $y_{T-1}$ than $x_{iT-1}$ was; thus unless differences in players' responses to new observations persist, the $x_{it}$ collapse mechanically on the current $y_t$; and if players differ only in their initial responses, $y_t$ remains forever at $y_0$, completely independent of the treatment variables.

5. Analyses of long-run equilibria of stochastic evolutionary dynamics have $\sigma^2_{\xi t} \equiv \varepsilon > 0$ (or $\sigma^2_{\xi t} \to \varepsilon > 0$) and $\alpha_t \equiv 0$ and $\beta_t \equiv \beta$ for all $t$ = 1, . . . , which yields (eventually) stationary transition probabilities and ergodic dynamics; the "long-run equilibrium" is defined as the support of the ergodic distribution as $\varepsilon \to 0$, which usually puts probability approaching one on one of the steady states; which one is determined by the relative difficulties of moving from alternative steady states to the basins of attraction of other steady states

● The analysis assumes little or no strategic sophistication and allows some strategic uncertainty.

● Such dynamics eventually lose the influence of all but recent history, with perpetually repeated "tunneling" across basins of attraction and no lock-in; but they normally imply a definite pattern of equilibrium selection, in a stochastic sense, which captures some of the effects of strategic uncertainty.

**Proposition 2**: In VHBB's 1990, 1991 games, the long-run equilibrium assigns probability one to the equilibrium with lowest (highest) effort whenever the order statistic is below (above) the median, and assigns positive probability to every equilibrium when the order statistic is the median. In each case the long-run equilibrium is independent of the number of players and the order statistic, as long as it remains below (or above) the median.

● The proof is a fairly simple mutation-counting argument like the one given in Robles *JET* 1997. Proposition 2 shows that a long-run equilibrium analysis discriminates among equilibria in ways qualitatively consistent with the variations across treatments VHBB observed, but otherwise eliminates most of the information about the effects of treatment variables.

**Estimation and further analysis**

6. Adaptive learning models have $\sigma^2_{\zeta 0} > 0$, $\sigma^2_{\zeta t} \to 0$, and $\alpha_t \to 0$ (or $\alpha_t \equiv 0$ for all $t$ = 1, . . .); the key difference between them and stochastic evolutionary dynamics is that $\sigma^2_{\zeta t} \to 0$ (rather than $\sigma^2_{\zeta t} \equiv \varepsilon > 0$ or $\sigma^2_{\zeta t} \to \varepsilon > 0$): this makes adaptive learning inherently nonstationary and nonergodic, while an analysis of long-run equilibria requires stationarity to make the dynamics ergodic (this stationarity must come either from new uncertainty each period or players continually entering and leaving the population)

● Adaptive learning models allow an extreme form of history-dependence, in which the dynamics lock in on a particular equilibrium in the stage game.

● A full analysis normally depends on the values of behavioral parameters; the model provides a framework in which to estimate them, using data from the experiments, and allowing different parameter values in each treatment.

● In the estimation, when the data allow it we impose the simplifying parameter restrictions $\beta_t \equiv \beta$ for all $t$ = 1, . . . ; we usually allow $\sigma^2_{\zeta t}$ to vary with $t$, though it is often useful to impose intertemporal restrictions of the form $\sigma^2_{\zeta t} = \sigma^2_{\zeta 1}/t^{\lambda}$ for $t$ = 1, . . . , where $\lambda \geq 0$.

● The estimated models give an adequate statistical summary of subjects' behavior, and generate dynamics and limiting outcomes in each treatment whose probability distributions closely resemble the empirical frequency distributions in the experiments.

● The estimates generally satisfy $\sigma^2_{\zeta 0} > 0$, $\sigma^2_{\zeta t} \to 0$, and $\alpha_t \to 0$, as one would expect as players learn to predict $y_t$ from common observations.

● Thus model that best describes behavior is an adaptive learning model with nonstationary transition probabilities and (because $\sigma^2_{\zeta t} \to 0$) possibly history-dependent dynamics.

● The rest of the analysis maintains assumptions general enough to include adaptive learning.

● Proposition 3 shows that unless $\sigma^2_{\zeta t} \to 0$ very slowly, the learning dynamics converge, with probability 1, to one of the symmetric equilibria of the coordination game; the variance condition needed here is just what one would guess from the strong law of large numbers

**Proposition 3**: Assume that the distributions of the $\zeta_{it}$ are truncated so that the $x_{it}$ remain in the interval $[\underline{x}, \bar{x}]$. Then if $0 < \beta \leq 1$ and $\sum_{s=0}^{\infty} \sigma^2_{\zeta s}$ is finite, $y_t$ and the $x_{it}$ converge, with probability 1, to a common, finite limit, which is an equilibrium of the stage game.

● The proof follows the martingale convergence arguments of Nevel'son and Has'minskii (1973, Theorem 2.7.3), using the stochastic Lyapunov function $V_t \equiv \sum_{i,j} (x_{it} - x_{jt})^2$.

● Given the conclusion of Proposition 3, the model's implications for equilibrium selection can be summarized by the prior probability distribution of the limiting equilibrium, which is normally nondegenerate due to the persistent effects of strategic uncertainty.

● The limiting outcome is determined by the cumulative drift before learning eliminates strategic uncertainty (faculty meeting example with varying quorum and group size).

● The form of the learning rules and the "evolutionary" structure of VHBB's designs allow a closed-form solution for $y_t$ and the $x_{it}$ as functions of the behavioral parameters, the treatment variables, and the shocks that represent strategic uncertainty:

**Proposition 4**: The unique solution of (3) and (4) is given, for all *i* and *t*, by

(5) $$x_{it} = \alpha_0 + \sum_{s=0}^{t-1} \beta f_s + z_{it}$$

and

(6) $$y_t = \alpha_0 + \sum_{s=0}^{t-1} \beta f_s + f_t,$$

where

(7) $$z_{it} \equiv \sum_{s=0}^{t} (1-\beta)^{t-s} \zeta_{is} \text{ and } f_t \equiv f(z_{1t},\ldots, z_{nt}).$$

● The proof is immediate by induction on *t* once the solution has been found. The solution is constructed uses the scaling property of *f(·)* and the linearity of the average adjustment rule in (2) to pass the common elements of the $x_{it}$ through *f(·)*.

● Proposition 4 shows how the outcome is built up period by period from the shocks that represent strategic uncertainty, whose effects persist indefinitely; the learning process is like a random walk in the aggregate, but with declining variances and nonzero drift.

● Persistence makes the limiting outcome depend on empirical behavioral parameters; this dependence is eliminated in other approaches only by ruling out either significant strategic uncertainty (as in equilibrium analyses) or its persistent effects (as in long-run equilibrium analyses of ergodic dynamics); paraphrase of quotation [about optimality, not equilibrium] in Stephen Jay Gould's *Wonderful Life*: "Equilibrium covers the tracks of history."

● Although players' beliefs and efforts become correlated via their responses to common observations of $y_t$, adopting an ex ante point of view and using Proposition 4's solution yields simple expressions for $Ex_{it}$ and $Ey_t$ in terms of the behavioral parameters, statistical parameters, and treatment variables.

● Let $\sigma_{zt}^2$ denote the common variance of the $z_{it}$. (7) implies that $\sigma_{zt}^2 \equiv \sum_{s=0}^{t}[(1-\beta)^{t-s}]^2 \sigma_{zs}^2$. Define $\mu_t \equiv Ef(z_{1t}/\sigma_{zt},\ldots,z_{nt}/\sigma_{zt})$. Because the $z_{it}/\sigma_{zt}$ are standardized, with mean 0 and variance 1, $\mu_t$ is completely determined by *n*, $f(\cdot)$, and the joint distribution of the $z_{it}/\sigma_{zt}$.

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

(11) $$Ex_{it} = \alpha_0 + \beta \sum_{s=0}^{t-1} \sigma_{zs}\mu_s$$

and

(12) $$Ey_t = \alpha_0 + \beta \sum_{s=0}^{t-1} \sigma_{zs}\mu_s + \sigma_{zt}\mu_t.$$

● The proof follows by taking expectations in (5) and (6), using (7), and noting that

(13) $$Ef(z_{1s},\ldots,z_{ns}) \equiv E[\sigma_{zs} f(z_{1s}/\sigma_{zs},\ldots,z_{ns}/\sigma_{zs})] \equiv \sigma_{zs}\mu_s.$$

● The proof exploits the fact that the shock terms in (5) and (6) are known functions of the $z_{it}$, which are ex ante i.i.d. across *i* with zero means for any given *t*.

● Proposition 5 shows how the drift that strategic uncertainty imparts to the dynamics depends on the variances that represent the dispersion of players' beliefs, behavioral parameters, and statistical parameters that reflect the influence of the treatment variables.

● Suppose, by way of approximation, that the $z_{it}/\sigma_{zt}$ are normal, so that $\mu_t \equiv \mu$. Suppose further that $\sum_{s=0}^{t}\sigma_{zs} \to S$ as $t \to \infty$. Then Propositions 3 and 5 imply that $Ey_t$ and $Ex_{it} \to \alpha_0 + \mu\beta S$. This shows how the mean coordination outcome is determined by the behavioral parameters $\alpha_0$ and $\beta$; the number of players and the order statistic, via $\mu$; and the initial dispersion of beliefs and the rate at which it is eliminated by learning, via $S$.

● By symmetry $\mu=0$ for VHBB's median and random-pairing minimum treatments, so there is no drift and $Ey_t$, $Ex_{it} \to \alpha_0$. The estimates of $\alpha_0$ were 4.30 in the random-pairing minimum treatment; 4.71 and 4.75 in median treatments Φ and Γ; and 6.26 in median treatment Ω (whose structure made the all-7 equilibrium more prominent). $\mu$ is negative (positive) for order statistics below (above) the median and strongly negative for the large-group minimum treatment, where, setting $n=15$, $\mu=-1.74$(Teichroew (1956, Table I). There the estimates of $\alpha_0$ and $\beta$ were 5.45 and 0.25, and $S$ (which is difficult to estimate for this treatment) appeared unlikely to be less than 10. Thus the approximate common limit of $Ey_t$, $Ex_{it}$ in the large-group minimum treatment, $\alpha_0 + \mu\beta S$, is < 1.10.

● For large initial levels of strategic uncertainty, declining gradually to zero over time as in our estimates, differences in drift across treatments make the prior probability distributions of the dynamics and limiting outcome vary with *n* and *f(·)* as their empirical frequency distributions varied in the experiments.

● Overall, the analysis yields the following conclusions:

a. Perfect history-dependence in 1991 median treatments is due to no drift and small variance; but convergence to initial median in 12 of 12 trials may overstate history-dependence: initial median "explains" 46-81% of variance of final median.

b. Lack of history-dependence in large-group minimum treatment is due to strong downward drift, which yields convergence to lower bound with very high probability; but convergence in 9 of 9 trials may understate the difficulty of coordination: in simulations it occurred in 500 of 500 trials.

c. Slow convergence, weak history-dependence, and lack of trend in the random-pairing minimum treatment are due to no drift and subjects' observation of only their current pair's minimum, which is a very noisy estimate of the population median that determined their best responses.

● The most important changes across treatments 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, the model treats the differences between these treatments primarily as changes in the order statistic (even though the former difference is "really" a change in group size and the latter also involves a change in group size, from 9 to 14-16). The above estimates suggest that each of these changes altered the drift of the process much more than the changes in behavioral parameters they induced.

● The analysis shows that similar results should be obtained in related environments, and yields qualitative comparative dynamics conclusions about the direct effects of changes in treatment variables, holding the behavioral parameters constant (they are not really constant, but estimates and analysis suggest that the former effects dominate):

d. Coordination is less efficient the lower the order statistic (the smaller the subsets of the population that can adversely affect the outcome), because small numbers of deviations are more likely than large numbers.

e. Coordination is to be less efficient in larger groups (holding the order statistic constant, measured from the bottom) because it requires coherence among more independent decisions (not up-down asymmetry!).

**Analysis of VHBB's 1993 result**

● The model extends earlier analyses to the 1993 auction environment, exploiting special features of design to analyze the dynamics of learning to play the two-stage game and how they interact with strategic uncertainty.

● Representing players' beliefs by optimal efforts allows us to describe decisions in both stages by a single beliefs variable as before, expectations about the order statistic.

● This allows us to describe dynamics of players' beliefs in terms of observable variables, in a way that permits estimation and analysis.

● The same approach explains efficiency-enhancing effect of auctions: the "order statistic" and "optimistic subjects" effects in VHBB's 1993 experiment can be approximated by adjusting the order statistic, so auctioning a 9-person median game in a group of 18 is like adjusting the median to the 75th percentile, as if all 18 played a game without auctions with payoffs and best responses determined by the 5th largest (the median of 9 largest) effort.

● Estimates and analysis suggest that there is also a "forward induction" effect, which contributes a roughly equal amount to the efficiency-enhancing effect of auctions.

● The analysis shows that the efficiency-enhancing effect should extend to related environments, but may not always yield full efficiency (e.g., moderately competitive auctions may not be enough to overcome the difficulty of coordination in large-group minimum games).

● The analysis also yields qualitative comparative dynamics conclusions, in which earlier results on changes in order statistic and group size remain valid, and a new result is obtained on the effects of increased competition for the right to play:

f. coordination is more efficient with more intense competition for the right to play, because it yields higher prices for a given level of dispersion in bidding strategies, and it increases the optimistic subjects effect.