Second-Year Advanced Microeconomics: Behavioural Economics

Behavioural Game Theory: Adaptive Learning, Hilary Term 2010

Vincent P. Crawford, University of Oxford

(with large debts to Colin Camerer and Teck-Hua Ho)

Last revised 22 December 2009

Introduction

Learning models describe how players adjust their decisions over time in response to experience with analogous games. (The learning process is usually modeled as repetition of a fixed "stage game," so that the analogies are perfect; but some recent work relaxes that assumption.)

The game is played either by a small group randomly selected from one or more populations—for example, random pairing to play a two-person game, with player roles filled either from the same or from identifiably separate populations—or sometimes by the entire population at once (as in Van Huyck, Cook, and Battalio's (1997 *JEBO*) "Continental Divide" game, discussed earlier).

Players view decisions in the stage game as the objects of choice, and the dynamics of their decisions are modeled directly, or indirectly in terms of their beliefs, with decisions modeled as best replies.

Players' decisions and roles are distinguished by commonly understood labels: the "language" in which they encode their experience, and in which any convention that emerges will be expressed.

The lectures start with a brief overview of the leading alternative models of learning:

- "Rational learning" models

- Deterministic evolutionary dynamics

- Long-run equilibria of stochastic evolutionary dynamics

- Adaptive learning models such as reinforcement learning, beliefs-based learning, and experience-weighted attraction ("EWA") learning, in which players' decisions in the stage game adjust in a direction that would increase payoffs, other things equal, given the current state of the system

These models' underlying assumptions differ in three main ways:

- Strategic sophistication—the extent to which players' beliefs and behavior reflect an analysis of their environment as a game, taking its structure and others' incentives into account

- Strategic uncertainty—the extent of players' heterogeneity, or uncertainty about others' decisions

- The extent to which the models seek to predict behavior entirely by theory, without empirical Information (which is closely related to the extent to which they eliminate history-dependence)

The lectures then continue by considering the theories in the light of a particular body of evidence, following Crawford, "Learning Dynamics…" in *The Evolution of Economic Diversity*, 2001.

Leading alternative models

A. Rational learning

In rational learning models, players' decisions in the stage game are determined by an equilibrium in the repeated game that describes the entire learning process, sometimes with a particular selection.

(A variant of rational learning sometimes encountered, not further discussed here, is the use of QRE in the stage game, with time-varying precision, to describe a learning process.)

B. Deterministic evolutionary dynamics

In deterministic evolutionary dynamics, a large population or populations of players repeatedly play a game, without or with distinguished roles.

Individual players normally play only pure actions, with payoffs determined by their own actions and the population action frequencies.

Players in a given player role are identical but for their actions.

In biology the law of motion of the population action frequencies is derived, usually with a functional form known as the replicator dynamics, from the assumption that players inherit their actions unchanged from their parents, who reproduce at rates proportional to their current payoffs.

In economics similar dynamics are derived from plausible assumptions about individual adjustment.

The usual goal is to identify the locally stable steady states of the dynamics.

If the dynamics converge, they converge to a steady state in which the actions that persist are optimal in the stage game, given the limiting action frequencies; thus, the limiting frequencies are in Nash equilibrium.

Even though players' actions are not rationally chosen—indeed, not even chosen—the population collectively "learns" the equilibrium as its frequencies evolve, with selection doing the work of rationality and strategic sophistication.

Because evolutionary models are useful templates for adaptive learning models, I give a detailed example of how this works.

Imagine a large population of men and women repeatedly and anonymously paired (with gender publicly observable in each pair, so can base their strategies on gender) to play Battle of the Sexes.

	Fights	**Ballet**
Fights	**2**, **1**	**0**, **0**
Ballet	**0**, **0**	**1**, **2**

Battle of the Sexes

Now draw a differential equation "phase diagram" with the population frequency of men playing Fights, *m*, on the horizontal axis and the frequency of women playing Fights, *w*, on the vertical axis. We will use this diagram to analyze the dynamics of simple learning rules.

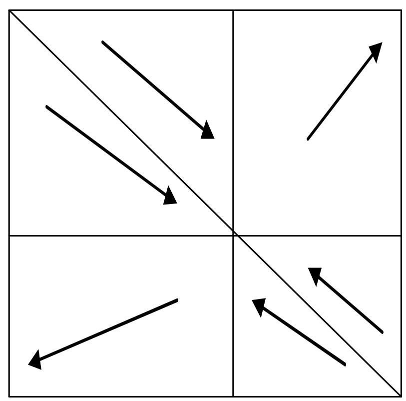

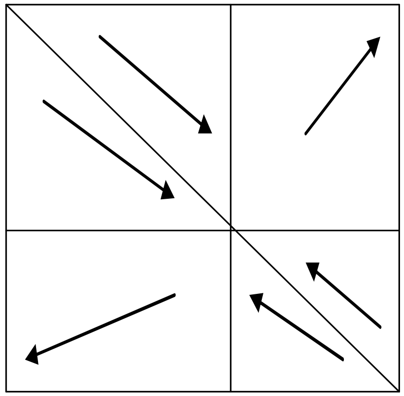

For men the expected payoff of Fights is higher than Ballet whenever $w > 1/3$ ($2w > 1 - w$). For women the expected payoff of Fights is higher whenever $m > 2/3$ ($m > 2(1 - m)$).

There are four regions: ($m > 2/3$, $w > 1/3$), ($m > 2/3$, $w < 1/3$), ($m < 2/3$, $w > 1/3$), ($m < 2/3$, $w < 1/3$).

For plausible learning rules, when ($m > 2/3$, $w > 1/3$), m and w rise. When ($m > 2/3$, $w < 1/3$), m falls and w rises. When ($m < 2/3$, $w > 1/3$), m rises and w falls. And when ($m < 2/3$, $w < 1/3$), m and w fall. When ($m > 2/3$, $w > 1/3$), $m \to 1$ and $w \to 1$; and when ($m < 2/3$, $w < 1/3$), $m \to 0$ and $w \to 0$. When ($m > 2/3$, $w < 1/3$) or ($m < 2/3$, $w > 1/3$), if (with symmetry) the initial condition is above the diagonal—$m + w > 1$—the system enters ($m > 2/3$, $w > 1/3$) and $m \to 1$ and $w \to 1$; if it's below the diagonal, the system enters ($m < 2/3$, $w < 1/3$) and $m \to 0$ and $w \to 0$.

In this setting the limiting outcome must be one of the two pure-strategy equilibria, in each of which all people follow a convention based on the commonly understood Fights versus Ballet labeling of their decisions. Which one they will follow is completely determined by whether the frequencies of initially arrogant men and wimpy women sum to more than half the population.

Now consider a large population repeatedly and anonymously paired to play the same kind of game, with two pure-strategy equilibria, one favored by one player and the other favored by the other; but with no observable labeling of players or decisions.

Players in this game can still use the payoffs to distinguish their strategies according to which one would yield them the more favorable outcome if their partner coordinated with it.

Follow the evolutionary game theory literature in calling these strategies Hawk (choose the strategy that would yield you the more favorable outcome if your partner coordinates with it, as Fights previously did for men and Ballet did for women) and Dove (choose the decision that would yield your partner the more favorable outcome if he coordinates with it).

With this redescription, in terms of labels that reflect the symmetry of men's and women's strategic positions, we can represent Battle of the Sexes symmetrically like this:

	Hawk	**Dove**
Hawk	**0**, **0**	**2**, **1**
Dove	**1**, **2**	**0**, **0**

Hawk-Dove Game

There are two equivalent ways to analyze the learning dynamics in this game.

The first is to recycle the phase diagram used to analyze Battle of the Sexes, but to impose the added restriction that the frequency of players playing Hawk must be equal in both player roles. This is just as if in Battle of the Sexes the frequency of men playing Fights, m in my notation, must be equal to the frequency of women playing Ballet, $1 - w$.

Because $m = 1 - w$ is equivalent to $m + w = 1$, this restriction limits the dynamics to the diagonal running from northwest to southeast in the previous two-dimensional phase diagram.

As the diagram suggests, the dynamics will now converge to the intersection of lines in the center, which represents the mixed-strategy equilibrium of the game at Pr{Hawk} = 2/3.

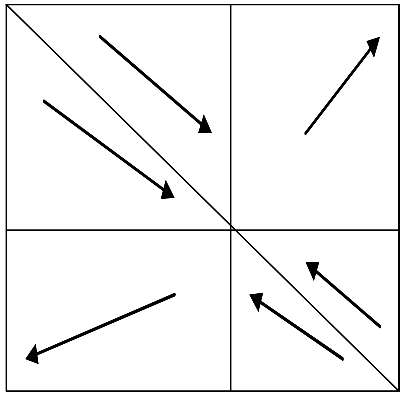

	Hawk	Dove
Hawk	0, 0	2, 1
Dove	1, 2	0, 0

Hawk-Dove Game

The second, less magical way to analyze the learning dynamics is to graph the expected payoffs of Hawk and Dove (in either player role) against the population frequency of Hawk.

This "builds in" the restriction that the frequency of players playing for their favorite equilibrium must be the same in both roles, and allows us to represent the dynamics in a one-dimensional phase diagram, with expected payoffs of Hawk and Dove on the vertical axis and population frequency of Hawk on the horizontal axis.

When the frequency of Hawk is low, Hawk has higher payoff than Dove, and vice versa. Thus the dynamics follow the arrows on the horizontal axis, converging to the frequency of Hawk where the payoff lines cross, which is 2/3, representing the mixed-strategy equilibrium.

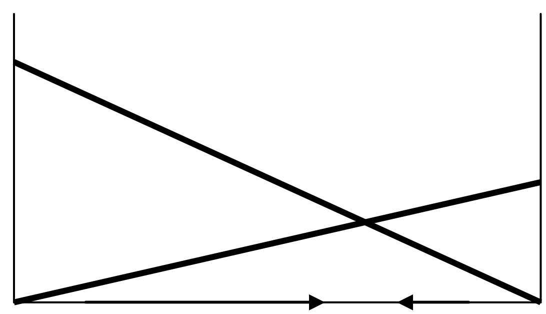

C. Long-run equilibria of stochastic evolutionary dynamics

Analyses of long-run equilibria of stochastic evolutionary dynamics (Kandori, Mailath, and Rob (1993 *Econometrica*) and Young (1993 *Econometrica*)) assume population interaction patterns like those in simple evolutionary game theory.

The state of the population is characterized by its current mix of strategies, and players' adjustments are assumed to move their strategies to or toward their best responses to the current state.

The main difference from deterministic evolutionary dynamics is that players' adjustments are subject to random "mutations," whose probability is constant over time and independent of the state.

D. Adaptive learning models

Reinforcement learning (from Camerer and Ho; see them for algebra)

One step up from evolutionary models in the cognitive sophistication agents are assumed to have are reinforcement approaches (also called stimulus-response or rote learning). Choice reinforcement assumes strategies are 'reinforced' by their previous payoffs, and the propensity to choose a strategy depends in some precise way on its stock of reinforcement. Reinforcement may also 'spill over' to strategies which appear to be similar to the chosen strategy (e.g., neighboring strategies, if strategies are rank-ordered) .

Reinforcement learners care only about the payoffs strategies yielded in the past, not about the history of play that created those payoffs. As a result, they will continue to play strictly dominated strategies if they began playing them early on and they performed well. Reinforcement learning is a reasonable theory for players with very imperfect reasoning ability (nonhuman animals pecking levers in the lab or foraging in the wild) or for human players who know nothing about the foregone payoffs from strategies they did not choose.[3]

Beliefs-based learning (from Camerer and Ho)

Belief-based models assume players update beliefs about what others will do, and use those beliefs to determine which strategies are best (by computing the expected payoffs to each strategy). A popular model is 'fictitious play'. In fictitious play, players keep track of the relative frequency with which another player has played each strategy in the past. These relative frequencies are beliefs about what that player will do in the upcoming period. Players then calculate expected payoffs for each strategy given these beliefs, and choose strategies with higher expected payoffs more frequently.

Fictitious play counts all previous observations equally. At the opposite extreme is Cournot best-response dynamics: Assume the strategy played most recently by others will be played again. Cournot dynamics weight the most recent past very heavily and dismiss or discard older experiences.

Experience-weighted attraction learning (from Camerer and Ho)

Experience-weighted attraction (EWA) learning was designed by myself and Teck Ho (1997) to combine the most appealing elements of reinforcement and weighted fictitious play in a hybrid or "gene-splice". The model adds three features to reinforcement and belief learning: An experience weight $N(t)$ (which increases in response to experience); a parameter which capture the rate at which attractions and the experience weight $N(t)$ are discounted; and δ, the weight players give to foregone payoffs from unchosen strategies. When parameters are restricted to have certain values, EWA reduces to a simple version of choice reinforcement in which only chosen strategies are reinforced. When parameters are restricted in a different way, so unchosen strategies are reinforced as strongly as chosen strategies are, and attractions are weighted averages rather than cumulations, EWA reduces exactly to weighted fictitious play.

EWA learning includes reinforcement learning and a class of weighted fictitious play beliefs-based models as special cases.

In EWA strategies have attractions which reflect prior predispositions, are updated based on payoff experience, and determine choice probabilities according to some rule (e.g., logit).

A key feature is a parameter δ which weights the strength of hypothetical reinforcement of strategies which were not chosen according to the payoff they would have yielded.

When $\delta = 0$ choice reinforcement results. When $\delta = 1$, levels of reinforcement of strategies are proportional to expected payoffs given beliefs based on past history.

Another key feature is the growth rates of attractions.

The EWA model controls the growth rates by two decay parameters, φ and ρ, which depreciate attractions and amount of experience separately.

When $\varphi = \rho$ belief-based models result; when $\rho = 0$ choice reinforcement results.

Estimates of δ are generally around .50, φ around 1, and ρ varies from 0 to φ.

EWA combines the best features of choice reinforcement and beliefs-based models, allowing attractions to begin and grow flexibly as choice reinforcement does, but reinforcing unchosen strategies substantially as belief-based models implicitly do.

Evidence: Convergence and equilibrium selection via learning in Van Huyck, Battalio, and Beil's (1990 *AER*, 1991 *QJE*, 1993 *GEB*) coordination experiments

VHBB's 1990 and 1991 experimental designs

Repeated play of player-role-symmetric coordination games in populations of subjects, interacting all at once ("large groups") or in pairs drawn randomly ("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 (minimum games were played either by the entire population of 14-16 or by random pairs, median games were played by the entire population of 9).

(Here I focus on the treatment with repeated random pairing (C_d), with only brief mention of the results for the treatment with initially random but thereafter fixed pairing (C_f).)

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 subject populations were large enough that subjects treated own influences on order statistic as negligible (the smallest "large" number in behavioral game theory is around four or five).

Payoff Table Γ

		Median value of X chosen						
		7	6	5	4	3	2	1
Your choice of X	7	1.30	1.15	0.90	0.55	0.10	−0.45	−1.10
	6	1.25	1.20	1.05	0.80	0.45	0.00	−0.55
	5	1.10	1.15	1.10	0.95	0.70	0.35	−0.10
	4	0.85	1.00	1.05	1.00	0.85	0.60	0.25
	3	0.50	0.75	0.90	0.95	0.90	0.75	0.50
	2	0.05	0.40	0.65	0.80	0.85	0.80	0.65
	1	−0.50	−0.05	0.30	0.55	0.70	0.75	0.70

Payoff Table A

		Smallest Value of X Chosen						
		7	6	5	4	3	2	1
Your Choice of X	7	1.30	1.10	0.90	0.70	0.50	0.30	0.10
	6	–	1.20	1.00	0.80	0.60	0.40	0.20
	5	–	–	1.10	0.90	0.70	0.50	0.30
	4	–	–	–	1.00	0.80	0.60	0.40
	3	–	–	–	–	0.90	0.70	0.50
	2	–	–	–	–	–	0.80	0.60
	1	–	–	–	–	–	–	0.70

VHBB's Leading Median and Minimum Payoff Tables

The random-pairing and large-group minimum games are larger versions of two-effort Stag Hunts.

Other Player

	Stag	Rabbit
Stag	2, 2	0, 1
Rabbit	1, 0	1, 1

Two-Person Stag Hunt

All Other Players

	All-Stag	Not All-Stag
Stag	2	0
Rabbit	1	1

***n*-Person Stag Hunt**

The stage games all have seven strict, symmetric, Pareto-ranked equilibria, with players' best responses an order statistic of the population effort distribution.

The games are like a 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.

Intuitively, efficient coordination is more difficult, the larger the quorum or the larger the group, other things equal; but traditional equilibrium analysis and refinements don't fully reflect this.

VHBB's 1990 and 1991 results

The five leading treatments all evoked similar initial responses (table from Crawford (1991 *GEB*)).

TABLE I

		Minimum treatment				
		A (%)	B (%)	A' (%)	C_d (%)	C_f (%)
Subject's initial effort	7	33 (31)	76 (84)	23 (25)	11 (37)	13 (42)
	6	10 (9)	1 (1)	1 (1)	1 (3)	0 (0)
	5	34 (32)	2 (2)	2 (2)	2 (7)	6 (19)
	4	18 (17)	5 (5)	7 (8)	5 (17)	2 (6)
	3	5 (5)	1 (1)	7 (8)	3 (10)	1 (3)
	2	5 (5)	1 (1)	17 (19)	1 (3)	1 (3)
	1	2 (2)	5 (5)	34 (37)	7 (23)	8 (26)
Totals		107 (101)	91 (99)	91 (100)	30 (100)	31 (99)

		Median treatment		
		Γ, Γdm (%)	Ω (%)	Φ (%)
Subject's initial effort	7	8 (15)	14 (52)	2 (7)
	6	4 (7)	1 (4)	3 (11)
	5	15 (28)	9 (33)	9 (33)
	4	19 (35)	3 (11)	11 (41)
	3	8 (15)	0 (0)	2 (7)
	2	0 (0)	0 (0)	0 (0)
	1	0 (0)	0 (0)	0 (0)
Totals		54 (100)	27 (100)	27 (99)

Inexperienced subjects' initial strategic thinking doesn't react strongly to order statistic or group size.

Thus the strong treatment effects in subsequent outcomes are due to the dynamics of learning.

Subjects almost always converged to some equilibrium.

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

- 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)

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

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

Comparing the first two reveals an "order statistic" or "robustness" effect, with coordination less efficient the smaller the groups that can disrupt desirable outcomes.

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

TABLE III
Median Choice for the First Ten Periods of All Experiments

	Period									
Treatment	1	2	3	4	5	6	7	8	9	10
Gamma										
Exp. 1	4	4	4	4	4	4	4*	4	4*	4*
Exp. 2	5	5	5	5	5	5	5	5	5	5
Exp. 3	5	5	5	5	5	5	5	5	5	5*
Gammadm										
Exp. 4	4	4	4	4	4	4*	4*	4*	4*	4*
Exp. 5	4	4	4	4*	4*	4*	4*	4*	4*	4*
Exp. 6	5	5	5	5	5	5	5	5*	5*	5*
Omega										
Exp. 7	7	7	7	7*	7*	7*	7*	7*	7*	7*
Exp. 8	5	5	5	5	5*	5*	5*	5*	5*	5*
Exp. 9	7	7	7*	7*	7*	7*	7*	7*	7*	7*
Phi										
Exp. 10	4	4	4	4	4*	4*	4*	4*	4*	4*
Exp. 11	5	5	5	5*	5*	5*	5*	5*	5*	5*
Exp. 12	5	6	5	5*	5*	5*	5*	5*	5*	5*

Notes. Exp. = experiment. * = indicates a mutual best response outcome.

TABLE 2—EXPERIMENTAL RESULTS FOR TREATMENT *A*

	Period									
	1	2	3	4	5	6	7	8	9	10
Experiment 1										
No. of 7's	8	1	1	0	0	0	0	0	0	1
No. of 6's	3	2	1	0	0	0	0	0	0	0
No. of 5's	2	3	2	1	0	0	1	0	0	0
No. of 4's	1	6	5	4	1	1	1	0	0	0
No. of 3's	1	2	5	5	4	1	1	1	0	1
No. of 2's	1	2	2	4	8	7	8	6	4	1
No. of 1's	0	0	0	2	3	7	5	9	12	13
Minimum	2	2	2	1	1	1	1	1	1	1
Experiment 2										
No. of 7's	4	0	1	0	0	0	0	0	0	1
No. of 6's	1	0	1	0	0	1	0	0	0	0
No. of 5's	3	3	2	1	0	0	1	1	0	1
No. of 4's	4	6	2	3	3	0	0	0	0	0
No. of 3's	1	4	2	5	0	1	1	0	1	0
No. of 2's	3	2	6	5	5	9	3	4	3	1
No. of 1's	0	1	2	2	8	5	11	11	12	13
Minimum	2	1	1	1	1	1	1	1	1	1
Experiment 3										
No. of 7's	4	4	1	0	1	1	1	0	0	2
No. of 6's	2	0	2	0	0	0	0	0	0	0
No. of 5's	5	6	1	1	1	0	0	0	0	0
No. of 4's	3	3	2	1	2	1	0	0	0	1
No. of 3's	0	0	7	6	0	2	3	0	0	0
No. of 2's	0	1	1	4	5	3	6	3	2	2
No. of 1's	0	0	0	2	5	7	4	11	12	9
Minimum	4	2	2	1	1	1	1	1	1	1
Experiment 4										
No. of 7's	6	0	1	1	0	0	1	0	0	0
No. of 6's	0	6	2	0	0	1	0	0	0	0
No. of 5's	8	5	5	5	0	1	0	0	0	0
No. of 4's	1	1	4	6	7	1	2	1	1	0
No. of 3's	0	2	3	2	4	3	2	2	1	0
No. of 2's	0	1	0	0	2	3	7	4	2	2
No. of 1's	0	0	0	1	2	6	3	8	11	13
Minimum	4	2	3	1	1	1	1	1	1	1

TABLE 2—EXPERIMENTAL RESULTS FOR TREATMENT *A*, Continued

	Period									
	1	2	3	4	5	6	7	8	9	10
Experiment 5										
No. of 7's	2	2	3	1	1	1	1	0	0	0
No. of 6's	1	3	1	0	0	0	0	0	0	0
No. of 5's	9	3	0	4	1	0	2	0	0	0
No. of 4's	3	4	6	2	1	2	0	2	1	1
No. of 3's	1	2	2	4	6	0	0	0	0	1
No. of 2's	0	2	2	3	4	6	5	2	5	3
No. of 1's	0	0	2	2	3	7	8	12	10	11
Minimum	3	2	1	1	1	1	1	1	1	1
Experiment 6										
No. of 7's	5	3	1	1	1	1	2	2	2	3
No. of 6's	2	0	0	0	1	0	0	0	0	0
No. of 5's	5	1	0	0	0	1	0	0	0	0
No. of 4's	2	3	4	0	0	0	0	0	0	0
No. of 3's	1	5	4	2	2	2	1	0	2	0
No. of 2's	0	2	4	5	3	3	6	4	5	5
No. of 1's	1	2	3	8	9	9	7	10	7	8
Minimum	1	1	1	1	1	1	1	1	1	1
Experiment 7										
No. of 7's	4	3	1	1	1	1	1	1	1	1
No. of 6's	1	0	0	0	0	0	0	0	0	0
No. of 5's	2	3	0	0	0	0	0	0	0	0
No. of 4's	4	0	1	2	1	0	0	0	0	0
No. of 3's	1	3	2	1	1	0	0	0	0	0
No. of 2's	1	3	2	2	4	4	4	4	5	3
No. of 1's	1	2	8	8	7	9	9	9	8	10
Minimum	1	1	1	1	1	1	1	1	1	1

TABLE 5—DISTRIBUTION OF ACTIONS FOR TREATMENT *C*: RANDOM PAIRINGS

	Period				
	21	22	23	24	25
Experiment 6					
No. of 7's	5	5	4	10	8
No. of 6's	0	1	3	0	0
No. of 5's	2	5	3	3	4
No. of 4's	3	1	1	1	1
No. of 3's	1	1	1	0	0
No. of 2's	1	1	2	2	2
No. of 1's	4	2	2	0	1
Experiment 7					
No. of 7's	–	–	6	5	5
No. of 6's	–	–	1	0	1
No. of 5's	–	–	0	3	0
No. of 4's	–	–	2	1	4
No. of 3's	–	–	2	0	0
No. of 2's	–	–	0	0	1
No. of 1's	–	–	3	5	3

Aside

In case you are wondering what happened with fixed rather than random pairing, here are the results.

There is clear evidence of "strategic teaching" (Camerer et al. (2002 *JET*), with 12 out of 14 pairs somehow "teaching" their way to the most efficient equilibrium.

Most subjects seemed to understand that strategic teaching is pointless with random pairing, because it's costly but others reap almost all of the benefits. But they used it effectively with fixed pairing.

Outcomes cannot be modeled taking stage-game strategies as the objects of choice, because teaching looks beyond current payoffs.

But it's not clear how to model them taking repeated-game strategies as the objects of choice either.

TABLE 4—EXPERIMENTAL RESULTS FOR TREATMENT C: FIXED PAIRINGS

	Period						
	21	22	23	24	25	16	27
Experiment 5							
Pair 1							
Subject 1	7	7	7	7	7	7	7
Subject 16	7	7	7	7	7	7	7
Minimum	7*	7*	7*	7*	7*	7*	7*
Pair 2							
Subject 2	7	2	7	7	7	7	7
Subject 15	1	7	3	6	7	7	7
Minimum	1	2	7	7	7	7	7
Pair 3							
Subject 3	1	1	1	1	1	1	1
Subject 14	1	1	7	1	1	1	7
Minimum	1*	1*	1	1*	1*	1*	1
Pair 4							
Subject 4	1	7	7	7	7	7	7
Subject 13	7	2	5	7	7	7	7
Minimum	1	2	5	7*	7*	7*	7*
Pair 5							
Subject 5	1	7	4	7	7	7	7
Subject 12	1	4	7	7	7	7	7
Minimum	1	4	4	7*	7*	7*	7*
Pair 6							
Subject 6	5	7	7	7	7	7	7
Subject 11	7	7	7	7	7	7	7
Minimum	5	7*	7*	7*	7*	7*	7*
Pair 7							
Subject 7	1	7	6	7	7	7	7
Subject 10	5	3	6	7	7	7	7
Minimum	1	3	6*	7*	7*	7*	7*

Pair 8							
Subject 8	7	6	6	7	7	7	7
Subject 9	3	5	7	7	7	7	7
Minimum	3	5	6	7*	7*	7*	7*
Experiment 6							
Pair 1							
Subject 2	7	7	4	5	6	6	7
Subject 15	2	3	6	6	7	7	7
Minimum	2	3	4	5	6	6	7*
Pair 2							
Subject 3	5	7	7	7	7	7	7
Subject 14	7	7	7	7	7	7	7
Minimum	5	7*	7*	7*	7*	7*	7*
Pair 3							
Subject 4	1	1	1	1	4	4	1
Subject 13	7	1	1	3	1	1	2
Minimum	1	1*	1*	1	1	1	1
Pair 4							
Subject 5	5	7	7	7	7	7	7
Subject 12	7	7	7	7	7	7	7
Minimum	5	7*	7*	7*	7*	7*	7*

TABLE 4—FIXED PAIRINGS, Continued

	Period						
	21	22	23	24	25	26	27
Pair 5							
Subject 6	4	5	7	7	7	7	7
Subject 11	4	5	7	7	7	7	7
Minimum	4*	5*	7*	7*	7*	7*	7*
Pair 6							
Subject 7	5	7	7	7	7	7	7
Subject 10	5	7	7	7	7	7	7
Minimum	5*	7*	7*	7*	7*	7*	7*

* – Denotes a mutual best-response outcome.

End of aside

VHBB's 1993 design and results

VHBB's (1993 *GEB*) design was the same as their 1991 design, with repeated play of one of the 1991 median games, but with the right to play auctioned each period to the highest 9 bidders in a population of 18 (an English clock auction, with the same price paid by all winning bidders).

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 bid the payoff of some equilibrium of the median game and play that equilibrium, unless others bid differently.

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: strong, precise selection among a wide range of equilibria.

Auctioning the right to play had a strong efficiency-enhancing effect via focusing subjects' beliefs on more efficient ways to coordinate—a new and potentially important mechanism by which competition promotes efficiency.

TABLE V

Distribution of Actions for Game Γ(9): EC Auction

	Period 6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Exp. 10															
Price	1.24	1.24	1.28	1.29	1.30	1.30	1.30	1.30	1.30	1.30	—	—	—	—	—
Undom. actions	≥6	≥6	7	7	7	7	7	7	7	7	—	—	—	—	—
# of 7s	7	8	9	8	9	9	9	9	9	9	—	—	—	—	—
# of 6s	2	1	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 5s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 4s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 3s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 2s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 1s	0	0	0	1	0	0	0	0	0	0	—	—	—	—	—
Median	7	7	7*	7	7*	7*	7*	7*	7*	7*	—	—	—	—	—
Exp. 11															
Price	1.00	1.20	1.29	1.30	1.29	1.30	1.29	1.29	1.30	1.30	—	—	—	—	—
Undom. actions	≥4	≥6	7	7	7	7	7	7	7	7	—	—	—	—	—
# of 7s	4	5	9	9	9	9	9	9	9	9	—	—	—	—	—
# of 6s	1	3	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 5s	2	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 4s	2	1	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 3s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 2s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 1s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
Median	6	7	7*	7*	7*	7*	7*	7*	7*	7*	—	—	—	—	—
Exp. 12															
Price	.95	1.04	1.08	1.10	1.15	1.20	1.25	1.25	1.30	1.30	1.30	1.30	1.30	1.30	1.30
Undom. actions	≥3	≥4	≥5	≥5	≥5	≥6	≥6	≥6	7	7	7	7	7	7	7
# of 7s	1	0	1	3	2	5	8	9	9	9	9	9	9	9	9
# of 6s	0	3	5	6	7	4	0	0	0	0	0	0	0	0	0
# of 5s	6	2	2	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	1	4	1	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	5	5	6	6	6	7	7	7	7*	7*	7*	7*	7*	7*	7*
Exp. 13															
Price	1.05	1.14	1.18	1.25	1.29	1.25	1.25	1.30	1.25	1.30	1.30	1.30	1.30	1.30	1.30
Undom. actions	≥4	≥5	≥6	≥6	7	≥6	≥6	7	≥6	7	7	7	7	7	7
# of 7s	2	2	4	6	9	9	9	9	9	9	9	9	9	9	9
# of 6s	1	6	5	3	0	0	0	0	0	0	0	0	0	0	0
# of 5s	3	1	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	5	6	6	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*
Exp. 14															
Price	1.05	1.15	1.27	1.25	1.25	1.30	1.30	1.25	1.30	1.30	1.25	1.30	1.30	1.30	1.30
Undom. actions	≥4	≥5	7	≥6	≥6	7	7	≥6	7	7	7	7	7	7	7
# of 7s	0	5	8	8	9	9	9	9	9	9	9	9	9	9	9
# of 6s	7	4	0	1	0	0	0	0	0	0	0	0	0	0	0
# of 5s	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	6	7	7	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*

Notes. * indicates mutual best response outcome. ___ Partitions actions into $FI(P)$ and the complement of $FI(P)$.

TABLE VI

Distribution of Actions for Game Γ(9): EC Auction and Experienced Subjects

	Period														
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
Exp. 7 ($M = 5$)															
Price	1.09	1.09	1.10	1.19	1.29	1.29	1.30	1.29	1.30	1.30	1.30	1.29	1.30	1.25	1.29
Undom. actions	≥5	≥5	≥5	≥6	7	7	7	7	7	7	7	7	7	≥6	7
# of 7s	0	0	2	5	9	9	9	9	9	9	9	9	9	5	9
# of 6s	2	1	5	4	0	0	0	0	0	0	0	0	0	[illegible]	0
# of 5s	6	8	2	0	0	0	0	0	0	0	0	0	0	[illegible]	0
# of 4s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	5	5	6	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*
Exp. 8 ($M = 5$)															
Price	1.09	1.25	1.28	1.29	1.30	1.29	1.30	1.30	1.29	1.30	1.29	1.30	1.29	1.30	1.30
Undom. actions	≥5	≥6	7	7	7	7	7	7	7	7	7	7	7	7	7
# of 7s	3	7	9	9	9	9	9	9	9	9	9	9	9	9	9
# of 6s	2	2	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 5s	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	6	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*
Exp. 9 ($M = 6$)															
Price	1.15	1.21	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29
Undom. actions	≥5	≥6	7	7	7	7	7	7	7	7	7	7	7	7	7
# of 7s	0	7	9	9	9	9	9	9	9	9	9	9	9	9	9
# of 6s	8	1	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 5s	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	6	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*

Notes. * indicates mutual best response outcome. ___ Partitions actions into $FT(P)$ and the complement of $FT(P)$.

Explaining VHBB's 1990 and 1991 results (omitting those for fixed pairing)

A. Rational learning

Rational learning is unhelpful explaining VHBB's 1990 and 1991 results: Any pattern of coordinated jumping from one pure-strategy equilibrium to another over time is a rational learning equilibrium.

(Similarly, QRE with time-varying precision yields little or no insight here.)

B. Deterministic evolutionary dynamics

VHBB's results can be mostly (but not entirely) understood via a simple evolutionary basin of attraction story proposed in Crawford (1991 *GEB*, 1995 *Econometrica*).

Deterministic evolutionary dynamics have two advantages over traditional equilibrium analyses (including rational learning models) for the purpose of explaining results like VHBB's:

Together with the dispersion of initial responses, the effect of the order statistic on the sizes of the basins of attraction begins to capture the interaction between strategic uncertainty and learning dynamics.

And the dynamics give a rudimentary account of history-dependent equilibrium selection, in which the population always converges to the equilibrium whose basin of attraction includes its initial state.

Imagine that there are only two efforts as in Stag Hunt, not seven:

Other Player

	Stag	**Rabbit**
Stag	2 2	1 0
Rabbit	0 1	1 1

Two-Person Stag Hunt

All Other Players

	All-Stag	**Not All-Stag**
Stag	2	0
Rabbit	1	1

***n*-Person Stag Hunt**

Graph the expected payoffs of high (Stag) and low (Rabbit) effort against the population frequency of high effort in the random pairing and large-group minimum games and the large-group median game.

In the large-group median game, the all-Stag and all-Rabbit equilibria are both locally stable.

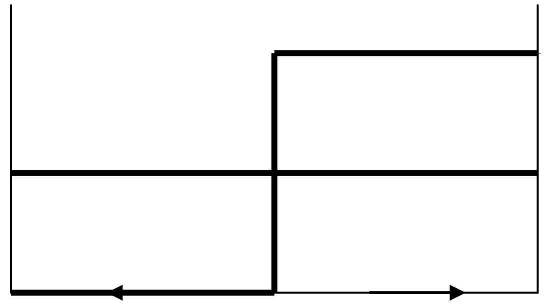

By symmetry, random shocks are neutral, equally likely to flip the population from all-Stag to all-Rabbit or vice versa.

With random initial conditions, the population would be equally likely to converge to all-Stag or all-Rabbit. If the initial conditions (strategic thinking) favor one equilibrium, then its probability of being selected is higher.

In the seven-effort version of the game that VHBB studied, if learning always makes subjects adjust their efforts toward the current value of the median, then the population converges to the median without changing it (a general property of order statistics like the median).

Even with random shocks, the median is just as likely to go up as it is to go down.

Either way, the learning dynamics have no up or down trend; and (given the dampening effect of the median on shocks) the population is very likely to “lock in” on the initial median, as it did in VHBB’s median experiments.

In the random-pairing minimum game, the all-Stag and all-Rabbit equilibria are again both locally stable.

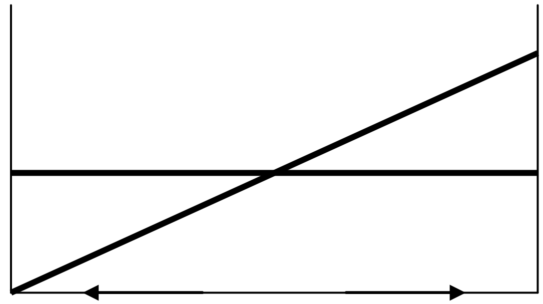

Random shocks are again neutral; and with random initial conditions, the population would be equally likely to converge to all-Stag or all-Rabbit.

Crawford (1995 *Econometrica*) shows that in the seven-effort version of this game that VHBB studied (i.e. for their payoffs), it’s actually optimal for a (risk-neutral) player to set his effort equal to his forecast of the median effort in the entire population.

Thus, just as in the large-group median game, the learning dynamics have no up or down trend and the population is likely to “lock in” on the initial median.

However, with random pairing a subject samples only a small fraction of the population effort distribution each period (his current partner’s effort is an estimate of the population median, but a very noisy one), so convergence will be much slower, as it was in VHBB’s experiments.

Table 5—Distribution of Actions for Treatment *C*: Random Pairings

	Period				
	21	22	23	24	25
Experiment 6					
No. of 7's	5	5	4	10	8
No. of 6's	0	1	3	0	0
No. of 5's	2	5	3	3	4
No. of 4's	3	1	1	1	1
No. of 3's	1	1	1	0	0
No. of 2's	1	1	2	2	2
No. of 1's	4	2	2	0	1
Experiment 7					
No. of 7's	–	–	6	5	5
No. of 6's	–	–	1	0	1
No. of 5's	–	–	0	3	0
No. of 4's	–	–	2	1	4
No. of 3's	–	–	2	0	0
No. of 2's	–	–	0	0	1
No. of 1's	–	–	3	5	3

Finally, in the large-group minimum game, the all-Rabbit equilibrium is locally stable but the all-Stag equilibrium is locally unstable. Starting from all-Stag, any shock, however small, will make the population converge to all-Rabbit.

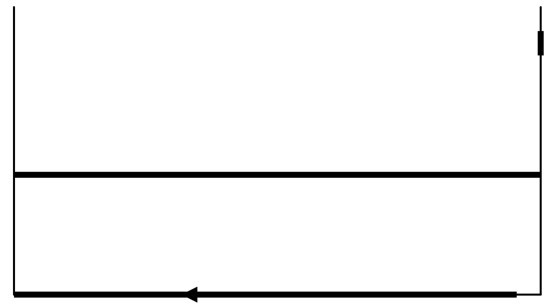

This makes the strong convergence to the equilibrium with lowest effort VHBB observed in the large-group minimum game plausible, but in this case the story is more complicated.

In the seven-effort large-group minimum game, if learning always made subjects adjust their efforts toward the current value of the minimum, then the population would converge monotonically to the initial minimum without ever changing it.

This result, formalized in Proposition 1 of Crawford, "Learning Dynamics…", is general across group sizes and order statistics in this class of games and evolutionary models.

However, in VHBB’s experiments the initial minimum was above one in five out of seven sessions, but it always converged quickly down to one. E.g.:

TABLE 2—EXPERIMENTAL RESULTS FOR TREATMENT *A*

	Period									
	1	2	3	4	5	6	7	8	9	10
Experiment 1										
No. of 7's	8	1	1	0	0	0	0	0	0	1
No. of 6's	3	2	1	0	0	0	0	0	0	0
No. of 5's	2	3	2	1	0	0	1	0	0	0
No. of 4's	1	6	5	4	1	1	1	0	0	0
No. of 3's	1	2	5	5	4	1	1	1	0	1
No. of 2's	1	2	2	4	8	7	8	6	4	1
No. of 1's	0	0	0	2	3	7	5	9	12	13
Minimum	2	2	2	1	1	1	1	1	1	1
Experiment 2										
No. of 7's	4	0	1	0	0	0	0	0	0	1
No. of 6's	1	0	1	0	0	1	0	0	0	0
No. of 5's	3	3	2	1	0	0	1	1	0	1
No. of 4's	4	6	2	3	3	0	0	0	0	0
No. of 3's	1	4	2	5	0	1	1	0	1	0
No. of 2's	3	2	6	5	5	9	3	4	3	1
No. of 1's	0	1	2	2	8	5	11	11	12	13
Minimum	2	1	1	1	1	1	1	1	1	1
Experiment 3										
No. of 7's	4	4	1	0	1	1	1	0	0	2
No. of 6's	2	0	2	0	0	0	0	0	0	0
No. of 5's	5	6	1	1	1	0	0	0	0	0
No. of 4's	3	3	2	1	2	1	0	0	0	1
No. of 3's	0	0	7	6	0	2	3	0	0	0
No. of 2's	0	1	1	4	5	3	6	3	2	2
No. of 1's	0	0	0	2	5	7	4	11	12	9
Minimum	4	2	2	1	1	1	1	1	1	1
Experiment 4										
No. of 7's	6	0	1	1	0	0	1	0	0	0
No. of 6's	0	6	2	0	0	1	0	0	0	0
No. of 5's	8	5	5	5	0	1	0	0	0	0
No. of 4's	1	1	4	6	7	1	2	1	1	0
No. of 3's	0	2	3	2	4	3	2	2	1	0
No. of 2's	0	1	0	0	2	3	7	4	2	2
No. of 1's	0	0	0	1	2	6	3	8	11	13
Minimum	4	2	3	1	1	1	1	1	1	1

Crawford (1995 *Econometrica*) shows that this happens because in the large-group minimum game, random shocks (which represent subjects' inability to perfectly predict others' adjustments) are not neutral as they were in the median game:

Instead they tend to make the minimum go down, to an extent that can be approximately quantified.

As in our intuition about the effect of a larger quorum or group suggests, the downward trend is stronger, the larger the group or the closer the order statistic (below the median) is to the minimum.

C. Long-run equilibria of stochastic evolutionary dynamics

Deterministic evolutionary dynamics may have many steady states, which in one-population models normally include the stage game's symmetric pure-strategy equilibria.

Which steady state the population converges to depends on the initial state and is difficult to predict.

But when the system is perpetually perturbed by random mutations, followed by adjustment to a best response to the current state, the limiting outcome may paradoxically become more predictable.

The dynamics are then ergodic, so that they converge to a distribution over states that is independent of history.

In the long run the process cycles perpetually among states, with their prior probabilities at any given time determined by the ergodic distribution.

This distribution depends on the probability of mutations, and is difficult to characterize in general.

But when the probability of mutations approaches zero, the ergodic distribution is concentrated around the steady states of the dynamics without mutations and approaches a limit that can be characterized by estimating the relative likelihoods of entering and exiting the steady states from the number of simultaneous mutations such changes require.

In VHBB's games, transitions between (symmetric pure-strategy) equilibria occur if and only if more players cross the order statistic from below than above, or vice versa.

Recall the large-group median Stag Hunt game, but contemplate alternative order statistics.

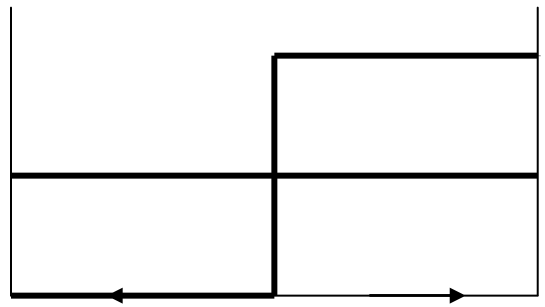

When the order statistic is below the median, the discontinuous drop in effort 2's expected payoff occurs in the left half of the horizontal axis.

When the basin of attraction of the low-effort equilibrium at the right end is larger than that of the high-effort equilibrium at the left, fewer simultaneous mutations are needed to go from the high-effort equilibrium to the edge of the basin of attraction of the low-effort equilibrium than vice versa.

A noninfinitesimal mutation probability therefore makes the probability of "tunneling" leftward across the boundary between basins of attraction lower than the probability of tunneling rightward, so that the ergodic distribution assigns higher probability to the low-effort equilibrium.

As the mutation probability approaches zero, the ratio of the two tunneling probabilities approaches infinity, and the probability of the low-effort equilibrium in the ergodic distribution approaches one.

This result, formalized in Proposition 2 of Crawford, "Learning Dynamics…" (see also Robles (1997 *JET*)), is general across group sizes and order statistics in this class of games and evolutionary models:

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 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.

Proposition 2 shows that analyses of long-run equilibria discriminate among strict equilibria and obtain unique predictions in most of VHBB's (1990, 1991) environments.

These predictions are obtained without modeling players' initial responses or using empirical information about behavior, by studying ergodic dynamics and passing to the limit as the mutation probability approaches zero.

However, although they distinguish between VHBB's large-group minimum treatment and their median and random-pairing minimum treatments in a simple way that is qualitatively consistent with the variations in observed outcomes, they are otherwise undiscriminating.

By limiting the effects of history, such analyses eliminate much of the information about the effects of changes in treatment variables an analysis of VHBB's results can provide.

D. Adaptive learning models

Crawford (1995 *Econometrica*), summarized in Crawford, "Learning Dynamics…", shows that the dynamics and limiting outcomes in VHBB's (1990 *AER*, 1991 *QJE*) games can be more fully understood via an adaptive learning model with heterogeneous beliefs.

The model assumes that players ignore their individual influences on the order statistic, learn to predict it, and independently choose their optimal efforts.

Learning is beliefs-based, which seems closest to what the evidence suggests here (although EWA may allow a somewhat better fit).

But learning is characterized in the style of the adaptive control literature, with players' beliefs represented by the optimal efforts they imply.

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.

The model is a Markov process with 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 the order statistic.

The key difference from stochastic evolutionary dynamics is that the heterogeneity of players’ beliefs, modeled as i.i.d. random perturbations about a common mean, converges to zero over time, rather than remaining with variance constant over time.

This makes adaptive learning inherently nonstationary and nonergodic, allowing the extreme form of history-dependence seen in the data, 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.

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.

Unless the heterogeneity of beliefs is eliminated very slowly, the learning dynamics converge, with probability 1, to one of the symmetric equilibria of the coordination game.

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 players' behavior as functions of the behavioral parameters, the treatment variables, and the shocks that represent strategic uncertainty, which shows how the outcome is built up period by period from the shocks that represent strategic uncertainty, whose effects persist indefinitely.

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).

Paraphrase of quotation [about optimality, not equilibrium] in Stephen Jay Gould's *Wonderful Life*:

"Equilibrium covers the tracks of history."

Overall, the analysis yields the following conclusions:

- 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.

- 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.

- 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 analysis yields qualitative comparative dynamics conclusions about the direct effects of changes in treatment variables, holding the behavioral parameters constant:

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.

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!).

Explaining VHBB's 1993 results

Crawford and Broseta (1998 *AER*), following Crawford (1995 *Econometrica*), show that this effect can be understood as following from effects that formalize "order statistic," "optimistic subjects," and "forward induction" intuitions.

The optimistic subjects and order statistic effects together have approximately the same magnitude in VHBB's environment (where the right to play a nine-person median game was auctioned in a group of 18) as the order statistic effect in an 18-person coordination game without auctions in which payoffs and best responses are determined by the fifth highest (the median of the nine highest) of all 18 players' efforts.

Auctioning the right to play a 9-person median game in a group of 18 effectively turns the game into a "75^{th} percentile" game (0.75 = 13.5/18), whose order statistic effect contributes a large upward drift as Crawford's (1995) analysis suggests there would have been in such a game without auctions.

Crawford and Broseta's analysis attributes the other half of the efficiency-enhancing effect of auctions in VHBB's environment to a strong forward induction effect.

The analysis shows that 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.

This effect should extend to related environments, but may not always yield full efficiency.