Undergraduate Behavioural and Experimental Game Theory Lectures, Hilary Term 2010

Vincent P. Crawford, University of Oxford

Last revised 24 December 2009

Introduction

Much of behavioural and experimental game theory focuses on two main issues:

- Strategic thinking, the process by which players predict others' decisions and make their own decisions in initial responses to games without clear precedents; and

- Adaptive learning, the process by which players learn to predict others' decisions from past experience with analogous games.

I begin with an experimental example showing why it is important to understand both issues, and then continue by discussing thinking and learning in turn.

Example: How strategic thinking and learning interact to determine equilibrium selection via history-dependent learning

In Van Huyck, Cook, and Battalio's (1997 *JEBO*) experiment, seven subjects chose simultaneously and anonymously among efforts from 1 to 14, with each subject's payoff determined by his own effort and a summary statistic, the median, of all players' efforts.

After subjects chose their efforts, the group median was publicly announced, subjects chose new efforts, and the process continued.

The relation between a subject's effort, the median effort, and his payoff was publicly announced via a table as on the next slide.

The payoffs of a player's best responses to each possible median are highlighted in bold in the table as displayed here (but not as displayed to subjects).

The payoffs of the (symmetric, pure-strategy) equilibria "all–3" and "all–12" are highlighted in large bold.

Continental divide game payoffs

	Median Choice													
Your Choice	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	45	49	52	55	56	55	46	-59	-88	-105	-117	-127	-135	-142
2	**48**	53	58	62	65	66	61	-27	-52	-67	-77	-86	-92	-98
3	**48**	**54**	**60**	**66**	70	74	72	1	-20	-32	-41	-48	-53	-58
4	43	51	58	65	**71**	**77**	80	26	8	-2	-9	-14	-19	-22
5	35	44	52	60	69	**77**	**83**	46	32	25	19	15	12	10
6	23	33	42	52	62	72	82	62	53	47	43	41	39	38
7	7	18	28	40	51	64	78	75	69	66	64	63	62	62
8	-13	-1	11	23	37	51	69	83	81	80	80	80	81	82
9	-37	-24	-11	3	18	35	57	88	89	91	92	94	96	98
10	-65	-51	-37	-21	-4	15	40	**89**	**94**	98	101	104	107	110
11	-97	-82	-66	-49	-31	-9	20	85	**94**	**100**	105	110	114	119
12	-133	-117	-100	-82	-61	-37	-5	78	91	99	**106**	**112**	**118**	**123**
13	-173	-156	-137	-118	-96	-69	-33	67	83	94	103	110	117	**123**
14	-217	-198	-179	-158	-134	-105	-65	52	72	85	95	104	112	120

There were ten sessions, each with its own separate group.

Half the groups happened to have an initial median of eight or above, and half happened to have an initial median of seven or below.

(The experimenters probably chose the design to make the initial median vary this way, but this kind of variation is not uncommon.)

The results are graphed on the next slide:

The median-eight-or-above groups converged almost perfectly to the all–12 equilibrium.

By contrast, the median-seven-or-below groups converged almost perfectly to the all–3 equilibrium.

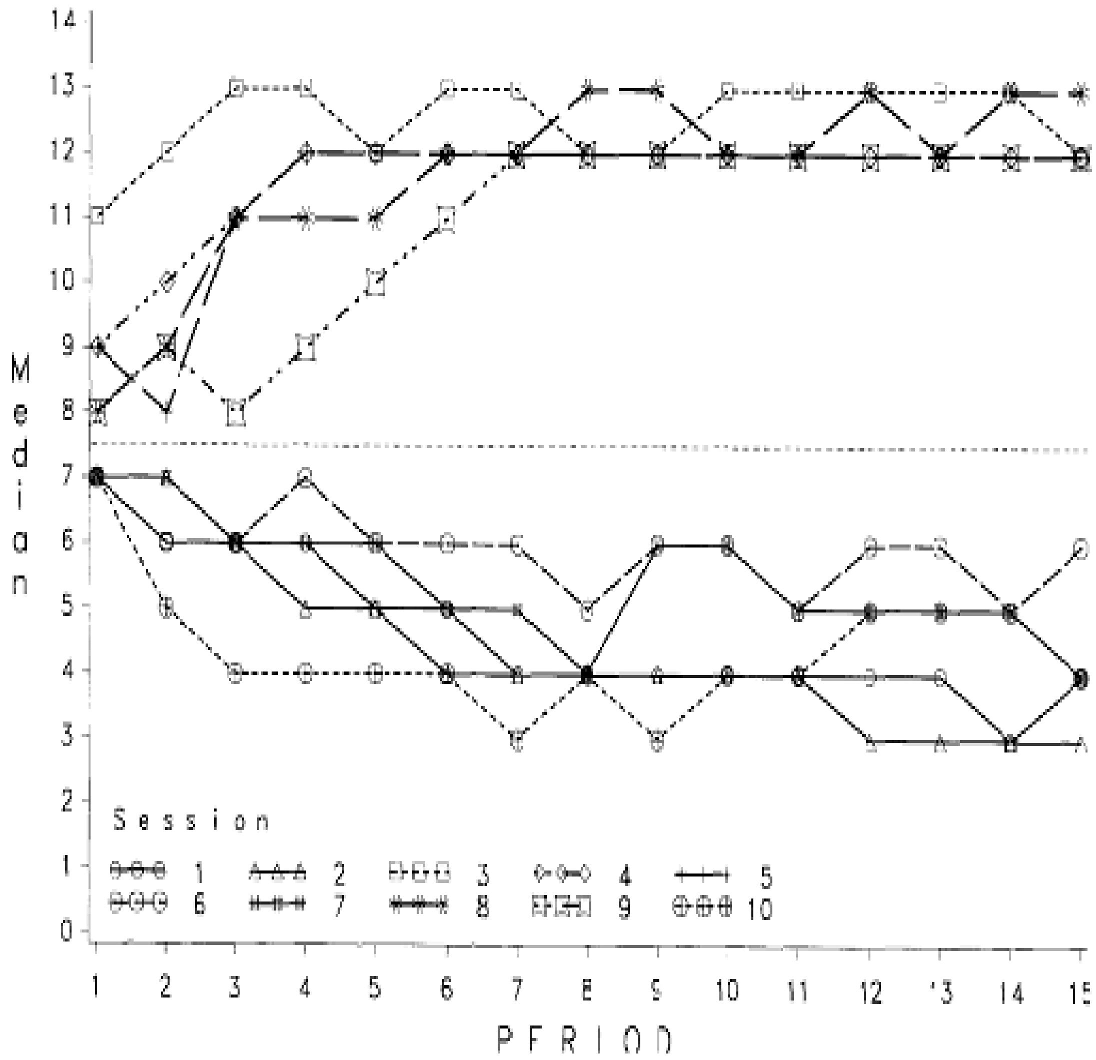

Fig. 3. Median choice in sessions 1 to 10 by period.

Van Huyck, Cook, and Battalio's Figure 3

Thus, it’s not enough to know that learning will eventually converge to some equilibrium, even if we are only interested in the final outcome.

Here we also need to know the prior probability distribution of the median initial response, which is determined by subjects' strategic thinking before they have any direct experience with the game.

In Van Huyck et al.’s experiment, the prior probability distribution of final outcomes is determined in a simple way by the prior probability distribution of the median initial response and learning rules that converge to the equilibrium whose basin of attraction—defined by myopic best responses—subjects’ initial responses fell into.

But in other settings predicting the prior probability distribution of final outcomes may require that we know more about the structure of subjects’ learning rules as well as their initial responses.

Strategic thinking

Strategic thinking is an essential part of human interaction, so much so that children must be *taught* to look both ways before crossing one-way streets.

(Once children develop enough "theory of mind" to distinguish others as independent decision makers, they seem to become instinctively overoptimistic about using rationality to predict others' decisions.)

Yet from a behavioral point of view, the importance of strategic thinking has been downplayed in economics and game theory.

Most applications of game theory in economics rely on Nash equilibrium.

But although equilibrium can be viewed as a model of strategic thinking, there are many potential applications of game theory for which it is not an adequate model of behavior.

Players' strategies will be in equilibrium if two conditions are satisfied:

- Players are rational (in the decision-theoretic sense of best responding to some beliefs).
- Players have the same beliefs about each other's strategies.

Assuming rationality for the sake of argument, there are two possible justifications for the assumption that players have the same beliefs:

- Thinking: If players have perfect models of each other's decisions, strategic thinking will lead them to have the same beliefs immediately, and so play an equilibrium even in their initial responses to a game.

  (Note that in this case the usual "as if" justification for equilibrium is unavailable: if players' models do not accurately reflect other players' cognition, equilibrium will predict their decisions accurately only by coincidence.)

- Learning: Even without perfect models, if players repeatedly play perfectly analogous games (and their interaction patterns do not foster repeated-game effects or strategic teaching), experience may eventually allow them to predict each others' decisions well enough to play an equilibrium (in the game that is repeated) in the limit.

In many applications of game theory, the theoretical conditions for learning to converge to equilibrium are approximately satisfied.

In such settings experimental evidence and field data tend to support assuming that players' steady-state strategies are in equilibrium.

If only long-run outcomes matter, and if equilibrium is unique or if there are multiple equilibria but equilibrium selection does not depend on the details of learning, such applications can safely rely entirely on equilibrium.

Because in such settings the cognitive requirements for learning to converge to equilibrium are mild, there is then no need to study strategic thinking.

However, many other applications involve games played without clear precedents, so that the learning justification for equilibrium is unavailable.

In other applications eventual convergence to equilibrium is assured, but initial as well as limiting outcomes matter (e.g. the FCC Spectrum auction).

In still other applications convergence is assured and only long-run outcomes matter, but the equilibrium is selected from multiple possibilities via history-dependent learning dynamics.

All such applications depend on reliably predicting initial responses to games, which may require a non-equilibrium model of strategic thinking.

Applications of game theory usually assume equilibrium even when its learning justification is unavailable.

This practice seems to be due to two factors:

- Fear that equilibrium is the only possible basis for analysis
  (rationalizability seldom yields predictions specific enough to be useful).

- Hope that equilibrium will still yield accurate predictions, on average.

But except in simple games, assuming equilibrium thinking in people's initial responses may be behaviorally far-fetched.

Even people who are capable of equilibrium thinking may doubt that others are capable, and therefore be unwilling to play their part of an equilibrium.

Moreover, there is a growing body of evidence—mostly experimental—that initial responses to novel or complex games often deviate systematically from equilibrium, especially if it requires thinking that is not straightforward.

Fortunately, the evidence also suggests that there are simple and tractable structural non-equilibrium models of strategic thinking that can explain a substantial fraction of people's deviations from equilibrium initial responses.

Those models allow equilibrium behavior, but do not assume equilibrium in all games.

Instead they assume that players follow strategic but non-equilibrium decision rules, which yield decisions that mimic equilibrium in simple games, but may deviate systematically in more complex games.

The models thereby provide a way to predict, in a given game, whether players' responses are likely to deviate from equilibrium, and if so, how.

Thus the hope that equilibrium yields predictions that are accurate on average is not well founded.

But neither is the fear that equilibrium is the only possible basis for analysis.

Modeling strategic thinking more accurately promises several benefits:

- It can establish the robustness of conclusions based on equilibrium in games where empirically reliable rules mimic equilibrium.

- It can challenge the conclusions of applications to games where equilibrium is implausible without learning.

- It can resolve empirical puzzles by explaining the deviations from equilibrium that some games evoke.

- It can also elucidate the structure of learning, where assumptions about cognition determine which analogies between current and previous games players recognize and also distinguish reinforcement from beliefs-based and more sophisticated rules.

“Folk game theory” quotations

I now give some folk game theory quotations to illustrate the need for models of strategic thinking, the issues successful models must address, and the range of potential applications.

Why study folk game theory instead of “real” game theory?

Folk game theory is only an imperfect reflection of traditional game theory, just as folk physics is an imperfect reflection of real physics.

But unlike folk physics, folk game theory has a direct and important influence on its observable counterpart, namely the part of behavioral game theory that concerns strategic thinking and initial responses to games.

I will argue below that the lessons regarding strategic thinking from folk game theory are largely confirmed by experiments designed to study strategic thinking in more conventional ways.

This correspondence is powerful evidence for a particular class of structural non-equilibrium models of strategic thinking.

Keynes's Beauty Contest:

"...professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. It is not a case of choosing those which, to the best of one's judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees."—John Maynard Keynes, *The General Theory of Employment*, *Interest, and Money*

(I suspect that the last sentence was Keynes's coy reference to himself.)

A simultaneous-move zero-sum *n*-person "outguessing" game, possibly with multiple equilibria. The key issue is anticipating others' strategic responses to a "landscape" of personal judgments about prettiness which is otherwise payoff-irrelevant. We will find that equilibrium alone is not very helpful in describing how people do this. The quotation suggests a thought process in which players "anchor" beliefs in instinctive reactions to the faces and then iterate best responses a finite number of times.

Kahneman's Entry Magic:

"…to a psychologist, it looks like magic."—Kahneman 1988, quoted in Camerer, Ho, and Chong (2004 *QJE*).

Here Kahneman refers to the fact that subjects in his market-entry experiments, structured like *n*-person Battle of the Sexes games, achieve better ex post coordination (number of entrants closer to market capacity) than in the natural symmetric mixed-strategy equilibrium benchmark.

(Thus Kahneman should have said "…to a *game theorist*, it looks like magic.")

The key issue here is breaking the symmetry of players' roles as required for efficient coordination. Equilibrium and refinements are not very helpful.

The same strategic issues arise in less abstractly framed, asymmetric field settings, exemplified by Roger Myerson's "Ware Medical Corporation" case (http://dss.ucsd.edu/~vcrawfor/Ware.htm):

A company is considering introducing a new product, which will be profitable only if its only competitor introduces a related product. The competitor's profits are determined qualitatively (not quantitatively) in the same way as the company's are. Both companies must decide, simultaneously and irreversibly, whether to begin development. In addition, there may be opportunities for commitment, signaling, and/or deceptive announcements….

Yushchenko:

“Any government wanting to kill an opponent…would not try it at a meeting with government officials.”—comment (quoted in the 2004 New York *Times*) on the poisoning of Ukrainian presidential candidate—now president—Viktor Yushchenko.

A simultaneous-move zero-sum two-person game with a unique mixed-strategy equilibrium. The players are a government assassin choosing one of several occasions at which to try to poison Yuschenko, only one of which is linked to the government; and an investigator who has the resources to check only one occasion.

Here the key issue is how players react to framing of decisions that is non-neutral but does not directly affect payoffs. Equilibrium in zero-sum two-person games leaves no room for such framing to affect outcomes, but people often react to it anyway.

The thinking reflected by the quotation is plainly strategic, but non-equilibrium: Any game theorist worth his salt would respond, “If that’s what people think, a meeting with government officials is exactly where *I* would try to poison Yushchenko.”

We will see that the quotation can be understood as a thought process in which a player anchors his beliefs in an instinctive reaction to the salience of the dinner with government officials and then iterates best responses a small number of times.

Lake Wobegon:

"…in Lake Wobegon, the correct answer is usually 'c'."—Garrison Keillor 1997 on multiple-choice tests (quoted in Attali and Bar-Hillel (2003 *Journal of Educational Measurement*)).

A simultaneous-move two-person zero-sum game with a unique mixed-strategy equilibrium. The players are a test designer deciding where to hide the correct answer and a clueless test-taker trying to guess the hiding place.

Again the key issue is how players react to the non-neutral framing, and the thinking reflected by the quotation is plainly strategic, but non-equilibrium.

Although there is nothing as uniquely salient as Yushchenko's dinner with government officials, psychologists like Christenfeld 1995 *Psychological Science* and Tversky (in Rubinstein, Tversky, and Heller 1996) think that with four possible answers, both the a and d end locations and location c are inherently salient (with the jury still out on which is more salient).

Again the quotation can be understood as a thought process in which a player anchors beliefs in an instinctive reaction to salience and iterates best responses a small number of times.

Common Features of the Quotations

- They all concern games played without completely clear precedents.

- They all reflect coherent, clearly identified models of strategic thinking.

- But the thinking is systematically different from equilibrium thinking.

- The thinking tends to start with beliefs anchored in an instinctive reaction to the game, and then to iterate best responses a small number of times.

  (In this respect the thinking resembles that in the "level-*k*" or "cognitive hierarchy" ("CH") models described below. The resemblance is not self-evident for Entry Magic, but as explained below, Camerer, Ho, and Chong (2004 *QJE*) explain Kahneman's results via a CH model.)

- The instinctive reactions follow different principles, each plausible in its setting, such as uniform randomness, salient labels, or truthfulness.

- Finite iteration of best responses is common across all settings, although the number of iterations may vary across individuals or even settings.

These common features are representative of folk game theory:

- One can also find quotations reflecting one or two steps of iterated (strict or weak) dominance in the normal form, or one or two steps of iterated (weak) dominance reflecting forward or backward induction in the extensive form.

- But it is difficult (counterexamples welcome) to find quotations involving more than one or two steps of iterated dominance.

- And it is at least as difficult (impossible? counterexamples welcome) to find quotations that illustrate the fixed-point reasoning that underlies equilibrium in games without dominance.

In Selten's 1998 *European Economic Review* words (but generalizing about the results of game experiments, not about folk game theory):

> "Basic concepts in game theory are often circular in the sense that they are based on definitions by implicit properties…. Boundedly rational strategic reasoning seems to avoid circular concepts. It directly results in a procedure by which a problem solution is found."

To paraphrase:

> "Real people don't use fixed-point reasoning to decide what to do."

This is not to say that with enough experience in a sufficiently stationary setting, learning can't make people converge to steady states that an *analyst* would need fixed-point reasoning to characterize.

Selten's point is simply that when equilibrium requires fixed-point reasoning, it may not be a good behavioral model of people's cognition.

Level-*k* models

Although the number of logically possible non-equilibrium model seems daunting, both folk game theory and experimental evidence support a particular class of models called level-*k* or cognitive hierarchy (CH) models.

Level-*k* models allow behavior to be heterogeneous, but assume that each player follows a rule drawn from a common distribution over a particular hierarchy of decision rules or *types* (as they are called in this literature; no relation to "types" as realizations of private information variables).

Type *Lk* anchors its beliefs in a nonstrategic *L0* type, which is meant to describe *Lk*'s model of others' instinctive reactions to the game.

The instinctive reactions may follow one of several principles depending on the setting, such as uniform randomness, salience, or truthfulness.

Lk then adjusts its beliefs via thought-experiments with iterated best responses: *L1* best responds to *L0*, *L2* to *L1*, and so on.

Like equilibrium players, *L1* and higher types are rational in that they choose best responses to beliefs, with perfect models of the game.

Lk's only departure from equilibrium is in replacing its perfect model of others' decisions with simplified models that avoid the complexity of equilibrium.

In applications it is usually assumed that *L1* and higher types make errors, which are often taken to be logit with estimated precision as in LQRE.

Thus the probability density of each type's decision is increasing in its expected payoff, evaluated using the type's model of others' decisions: *L2*, for example, makes errors whose distribution is sensitive to the payoff costs of deviations, evaluated assuming that other players' decisions are *L1*.

Unlike LQRE, *Lk* types do not respond to the noisiness of others' decisions.

Even so, the deterministic structure of a level-*k* model captures the sensitivity of players' deviations from equilibrium to out-of-equilibrium payoffs.

(Level-*k* models are thus structural alternatives to models like quantal response equilibrium, which treat deviations from equilibrium entirely as players' responses to others' errors.)

The population type frequencies are treated as behavioral parameters, to be estimated from the data or translated or extrapolated from previous analyses.

The estimated type distribution is typically fairly stable across games, with most weight on *L1*, *L2*, and perhaps *L3*.

The estimated frequency of the anchoring *L0* type is usually small.

Thus, *L0* "exists" mainly as *L1*'s model of others, *L2*'s model of *L1*'s model of others, and so on.

Low frequencies of *L0* are an important sign of health for a level-*k* model, in that high frequencies of *L0* would reduce the model to a parameterized distribution of responses, thus describing the data rather than explaining it.

Only when the strategic iteration of best responses plays a role can the model yield a useful explanation of the data.

Even though *L0* normally has a low frequency, its specification is the main issue in defining a level-*k* model and the key to its explanatory power.

As illustrated below, *L0* needs to be adapted to the setting, and there is an emerging consensus about how to do this in particular applications.

By contrast, the definition of *L1*, *L2*, and *L3* via iterated best responses allows a simple, reliable explanation of behavior across different settings.

Like equilibrium plus noise and QRE, level-*k* models are general models of strategic behavior, with small numbers of behavioral parameters.

Like CH models, discussed below, level-*k* models make point predictions that depend only on *L0* and the estimated type distribution.

L1 and higher types make undominated decisions, and *Lk* complies with *k* rounds of iterated dominance and *k*–rationalizability (thanks to Robert Östling of Stockholm University for clarifying this relationship).

Thus, a distribution of *Lk* types realistically concentrated on low levels of *k* mimics equilibrium in games that are dominance-solvable in a few rounds.

But such a distribution deviates systematically from equilibrium in some more complex games, in predictable ways.

These features allow level-*k* models to capture the sensitivity of deviations from equilibrium to out-of-equilibrium payoffs.

As a result, like LQRE, level-*k* (and CH) models often fit initial responses better than equilibrium plus noise.

Cognitive hierarchy ("CH") models

In Camerer, Ho, and Chong's (2004 *QJE*) cognitive hierarchy ("CH") model, a close relative of level-*k* models, *Lk* best responds not to *Lk-1* alone but to an estimated mixture of lower-level types; and the type frequencies are not unrestricted, but instead are treated as a parameterized Poisson distribution.

For an outside observer modeling behavior econometrically, this estimated- mixture specification seems more natural than the level-*k* specification.

But which specification better describes people's strategic thinking remains an empirical question (on which the jury is still not completely in).

A CH *L1* is the same as a level-*k* *L1*, but CH *L2* and higher types may differ.

A CH *L1* and higher types make undominated decisions, but unlike level-*k* types, but a CH *Lk* might not comply with *k* rounds of iterated dominance and *k*–rationalizability.

Unlike in a level-*k* model, in a CH model *L1* and higher types are usually assumed not to make errors.

Instead the uniformly random *L0*, which has positive frequency in the Poisson distribution, doubles as an error structure for *L1* and higher types.

A CH model makes point predictions that depend only on *L0* and the estimated Poisson parameter.

In some applications the Poisson constraint, imposed as a simplifying restriction, is not very restrictive and the CH model fits as well as a level-*k* model; but in others the Poisson constraint is strongly binding.

Experimental evidence

Level-*k* and CH models are now supported by a large body of experimental evidence on initial responses to games with various structures.

Here I focus on two representative experiments with normal-form games:

- Nagel's 1995 *AER* experiments, which were directly inspired by Keynes's Beauty Contest, and which provide a simple introduction to the evidence and the class of models that it suggests.

- Costa-Gomes and Crawford's (CGC) 2006 *AER* experiments, which use a much more powerful design to identify subjects' strategic thinking more precisely.

CGC's conclusions are fully consistent with the conclusions of other studies of initial responses to abstract normal-form games, just more precise.

With adjustments described below, CGC's conclusions are also consistent with those of the studies of the other kinds of games mentioned above.

Nagel's design and results

In Nagel's *n*-person guessing game design:

- 15-18 subjects simultaneously guessed between [0,100].

- The subject whose guess was closest to a target p (= 1/2 or 2/3, say), times the group average guess wins a prize, say $50.

- The structure was publicly announced.

If you have not already done so, please take a moment to decide what you would guess, in a group of non-game-theorists:

- if $p = 1/2$,

- if $p = 2/3$.

Nagel's games have a unique equilibrium, in which all players guess 0.

The games are dominance-solvable, so the equilibrium can be found by iteratively eliminating dominated guesses.

For example, if $p = 1/2$:

- It's dominated to guess more than 50 (because $1/2 \times 100 \leq 50$).

- Unless you think that other people will make dominated guesses, it's also dominated to guess more than 25 (because $1/2 \times 50 \leq 25$).

- And so on, down to 12.5, 6.25, 3.125, and eventually to 0.

The rationality-based argument for this “all–0” equilibrium is stronger than many equilibrium arguments, because it depends only on iterated knowledge of rationality, not on the assumption that players have the same beliefs.

However, even people who are rational are seldom certain that others are rational, or that others believe that others are rational.

Thus, they won’t (and shouldn’t) guess 0. But what do (should) they do?

Nagel's subjects played these games repeatedly, but we can view their initial guesses as responses to games played as if in isolation if they treated their influences on the future as negligible, which is plausible in groups of 15 to 18.

Nagel's subjects never played their equilibrium strategies initially, and their responses deviated systematically from equilibrium.

Instead there were spikes that suggest a distribution of discrete thinking "types," respecting 0 to 3 rounds of iterated dominance in each treatment (next slide).

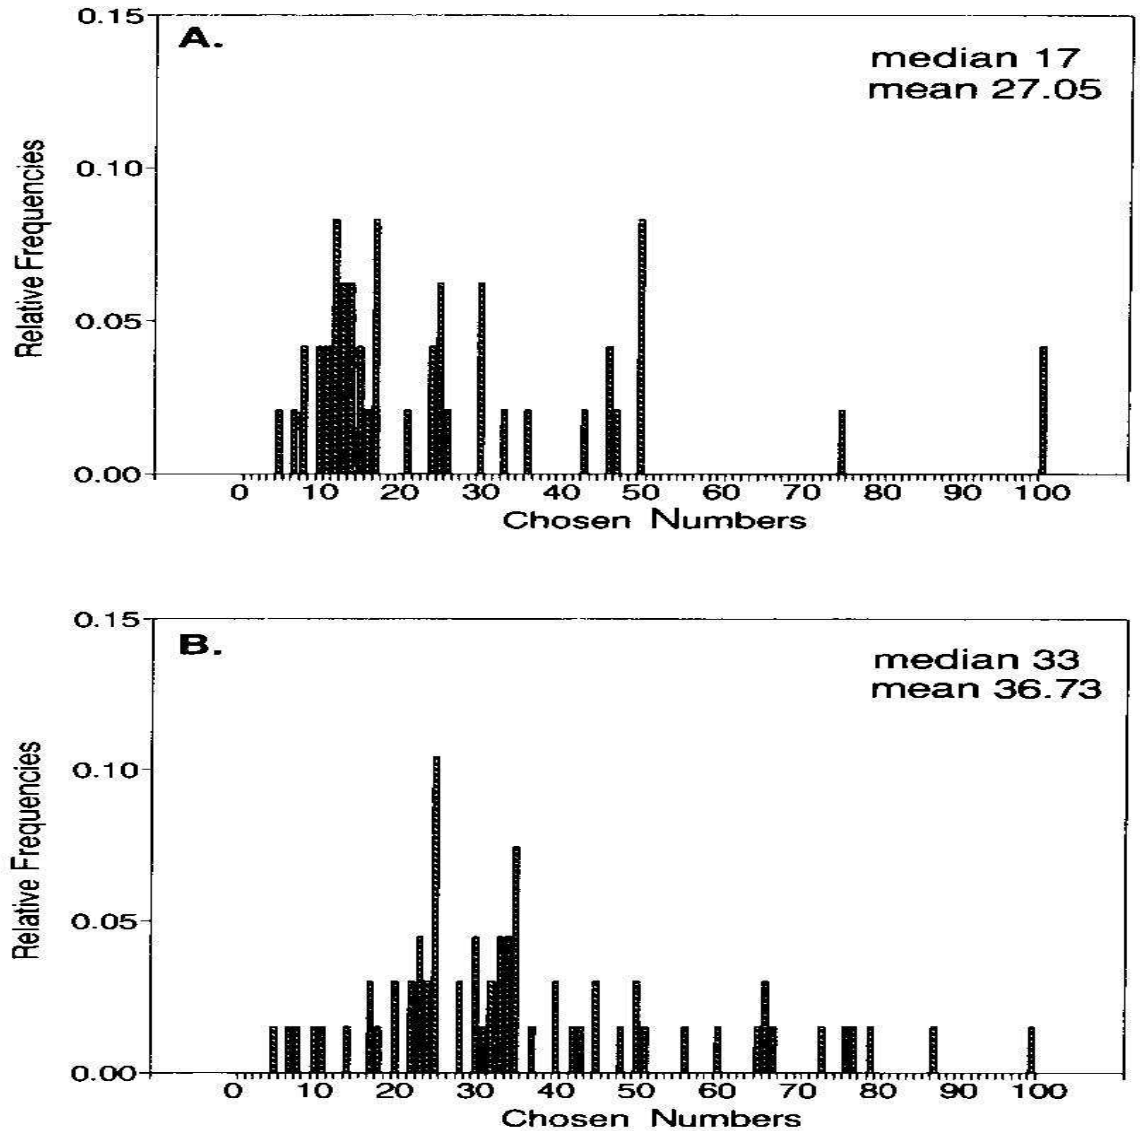

Part of Nagel's Figure 1: top of figure *p* = 1/2, bottom of figure *p* = 2/3.

The spikes' locations and how they vary across treatments are roughly consistent with two plausible interpretations:

- In one interpretation, called *Dk*, a player does *k* rounds of iterated dominance for some small number, *k* = 1 or 2, and then best responds to a uniform prior over other players' remaining strategies (thus "completing" *k*-rationalizability by adding a specific selection as discussed below).

- In another interpretation, "level-*k*" or "*Lk*," a player starts with a naïve prior *L0* over others' strategies reflecting people's instinctive reactions to the game, and then iterates best responses *k* times, with *k* = 1, 2, or 3.

In abstractly framed games like Nagel's, *L0* is usually taken to be a uniform random distribution, reflecting a player's understanding of the payoff function before he tries to model others' decisions. (In games without dominance this makes *Dk*, *k* = 1,2,… coincide with *L1*.)

(Although in these lectures I focus mainly on two-person games, in *n*-person games it matters whether *L0* is independent across players or correlated, and the limited evidence (HCW, Costa-Gomes, Crawford, and Iriberri 2009 *JEEA*) suggests that most people have highly correlated models of others. Here I take *L0* to model all others' average guess.)

In many games *Dk* and *Lk*+1 respond similarly to dominance, yielding *k*-rationalizable strategies. (The difference in indices is only a quirk of notation.)

With a uniform random *L0*, in Nagel's games *Dk*'s and *Lk*+1's guesses are perfectly confounded, both tracking the spikes in Nagel's data across her treatments (which had different subject groups):

- *Dk* guesses $([0+100p^k]/2)p$.
- *Lk*+1 guesses $[(0+100)/2]p^{k+1}$.

Either way, one aspect of the message is already clear: Subjects do not rely on indefinitely iterated dominance or indefinitely iterated best responses; instead their decisions respect *k*-rationalizability for at most small values of *k*.

Despite the lack of separation of *Dk*'s and *Lk*+1's guesses, many theorists interpret Nagel's results as evidence that subjects explicitly performed finitely iterated dominance, the way we teach students to solve such games.

In previous experiments, *Dk*'s and *Lk*+1's guesses were weakly separated, and the results are inconclusive on this point; but in CGC's experiments *Dk*'s and *Lk*+1's guesses are strongly separated, and we will see that the results very clearly favor *Lk* over *Dk* rules.

Costa-Gomes and Crawford's design and results

In CGC's design, subjects were randomly and anonymously paired to play a series of 16 different two-person guessing games, with no feedback.

The design suppresses learning and repeated-game effects in order to elicit subjects' initial responses, game by game, studying strategic thinking "uncontaminated" by learning.

("Eureka!" learning was possible, but it was tested for and found to be rare.)

The design combines the variation of games of Stahl and Wilson's 1995 *GEB* design with the large strategy spaces of Nagel's 1995 *AER* design.

This greatly enhances its power, and the profile of a subject's guesses in the 16 games forms a "fingerprint" that helps to identify his strategic thinking more precisely than is possible by observing his responses to a series of games with small strategy spaces or a single game with large strategy space.

In CGC's guessing games, each player has his own lower and upper limit, both strictly positive, implying finite dominance-solvability.

(Players are not actually required to guess between their limits. Instead guesses outside the limits are automatically adjusted up to the lower limit or down to the upper limit as necessary: a trick to enhance separation of information search implications, not important for this discussion.)

Each player also has his own target, and his payoff increases with the closeness of his guess to his target times the other's guess.

The targets and limits vary independently across players and games, with targets both less than one, both greater than one, or "mixed".

(In Nagel's and HCW's previous guessing experiments, the targets and limits were always the same for both players, and they varied at most across treatments with different subject groups.)

CGC's guessing games have essentially unique equilibria ("essentially" due to the automatic adjustment), determined (not always directly) by players' lower (upper) limits when the product of targets is less (greater) than one.

The discontinuity of the equilibrium correspondence when the product of targets equals one stress-tests equilibrium, which responds much more strongly to the product of the targets than alternative decision rules do; and enhances the separation of equilibrium from alternative rules.

(It also reveals other interesting patterns; see Crawford, "Look-ups as the Windows of the Strategic Soul".)

Consider a game in which players' targets are 0.7 and 1.5, the first player's limits are [300, 500], and the second's are [100, 900].

The product of targets is 1.05 > 1, and it can be shown that the equilibrium is therefore determined by players' upper limits. (When the product of targets is < 1, the equilibrium is determined by their lower limits in a similar way.)

In equilibrium the first player guesses his upper limit of 500, but the second player guesses 750 (= 500 × his target 1.5), below his upper limit of 900.

No guess is dominated for the first player, but any guess outside [450, 750] is dominated for the second player.

Given this, any guess outside [315, 500] is iteratively dominated for the first player.

Given this, any guess outside [472.5, 750] is dominated for the second player, and so on until the equilibrium at (500, 750) is reached after 22 rounds of iterated dominance.

Costa-Gomes and Crawford's data analysis

As suggested by previous work, CGC's data analysis assumed that each subject's guesses were determined, up to logit errors, by a single decision rule, or "type" as they are called in this literature (no relation to the use of "type" for the realization of a private information variable), in all 16 games.

This assumption was tested and found reasonable for almost all subjects.

Most of CGC's data analysis restricted attention to a list of behaviorally plausible types whose relevance was suggested by previous work:

- *L0*, *L1*, *L2*, and *L3*, with *L0* uniform random between a player's limits, *L1* best responding to *L0*, *L2* to *L1*, and so on.

- *D1* and *D2*, which does one round (respectively, two) of iterated dominance and then best responds to a uniform prior over its partner's remaining decisions (making a specific selection from *k*-rationalizable strategies).

- *Equilibrium*, which makes its equilibrium decisions.

  (Note that because CGC's games are all (finitely) dominance-solvable, traditional equilibrium refinements are not relevant in them.)

- *Sophisticated*, which best responds to the probability distributions of others' decisions, estimated from the observed frequencies.

  (*Sophisticated* is an ideal, included to learn if any subjects have an understanding of others' decisions that transcends mechanical rules.)

The restriction to this list was also tested and found to be a reasonable approximation to the support of subjects' decision rules.

CGC's large strategy spaces and the independent variation of targets and limits across games greatly enhance the separation of types' implications, to the point where many subjects' types can be precisely identified from their guessing "fingerprints":

Types' guesses in the 16 games, in (randomized) order played

	L1	***L2***	***L3***	***D1***	***D2***	***Eq.***	***Soph.***
1	**600**	**525**	**630**	**600**	**611.25**	**750**	**630**
2	**520**	**650**	**650**	**617.5**	**650**	**650**	**650**
3	**780**	**900**	**900**	**838.5**	**900**	**900**	**900**
4	**350**	**546**	**318.5**	**451.5**	**423.15**	**300**	**420**
5	**450**	**315**	**472.5**	**337.5**	**341.25**	**500**	**375**
6	**350**	**105**	**122.5**	**122.5**	**122.5**	**100**	**122**
7	**210**	**315**	**220.5**	**227.5**	**227.5**	**350**	**262**
8	**350**	**420**	**367.5**	**420**	**420**	**500**	**420**
9	**500**	**500**	**500**	**500**	**500**	**500**	**500**
10	**350**	**300**	**300**	**300**	**300**	**300**	**300**
11	**500**	**225**	**375**	**262.5**	**262.5**	**150**	**300**
12	**780**	**900**	**900**	**838.5**	**900**	**900**	**900**
13	**780**	**455**	**709.8**	**604.5**	**604.5**	**390**	**695**
14	**200**	**175**	**150**	**200**	**150**	**150**	**162**
15	**150**	**175**	**100**	**150**	**100**	**100**	**132**
16	**150**	**250**	**112.5**	**162.5**	**131.25**	**100**	**187**

Of the 88 subjects in CGC's main treatments, 43 made guesses that complied *exactly* (within 0.5) with one type's guesses in from 7 to 16 of the games (20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*).

For example, CGC's Figure 2 (next slide) shows the "fingerprints" of the 12 subjects whose guesses conformed most closely to *L2*'s; 72% of their guesses were exact *L2* guesses; only their deviations are shown.

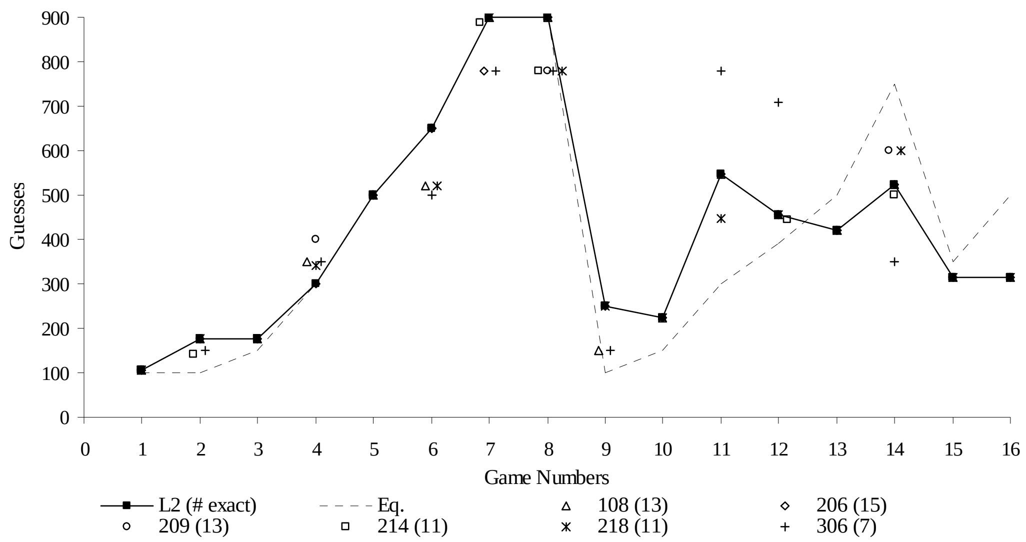

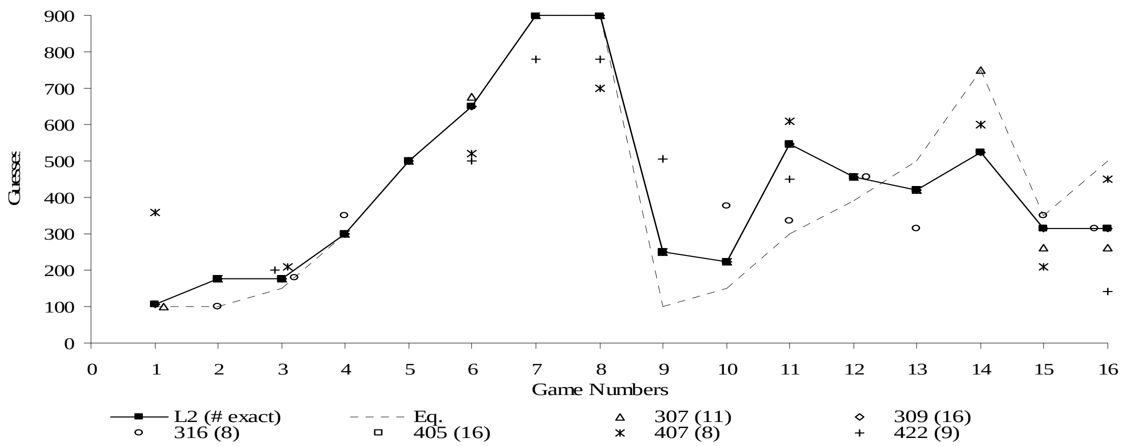

CGC's Figure 2. "Fingerprints" of 12 Apparent *L2* Subjects (Only deviations from *L2*'s guesses are shown.

(Of these subjects' 192 guesses, 138 (72%) were exact *L2* guesses.)

The size of CGC's strategy spaces, with 200 to 800 possible exact guesses in each of 16 different games, makes exact compliance powerful evidence for the type whose guesses are tracked: If a subject chooses 525, 650, 900 in games 1-3, intuitively and econometrically we already "know" he's an *L2*.

(By contrast, there are usually many possible reasons for choosing one of the strategies in a small matrix game; and even in Nagel's large strategy spaces, rules as cognitively disparate as *Dk* and *Lk*+1 yield identical decisions.)

Further, because CGC's definition of *L2* builds in risk-neutral, self-interested rationality, we also know that a subject's deviations from equilibrium are "caused" not by irrationality, risk aversion, altruism, spite, or confusion, but by his simplified model of others.

(Even so, doubts remain about the subjects with high exact compliance with *Equilibrium*, who appear to be following hybrid types that only mimic equilibrium in the games with targets both less than one or both greater than one; see Crawford, "Look-ups as the Windows of the Strategic Soul".)

That the level-*k* model is *directly* suggested by these subjects' data (rather than via data-fitting exercises) is an important advantage over alternatives.

CGC's other 45 subjects made guesses that conformed less closely to one of CGC's types, but econometric estimates of their types are concentrated on *L1*, *L2*, *L3*, and *Equilibrium*, in roughly the same proportions.

TABLE 1—SUMMARY OF BASELINE AND OB SUBJECTS' ESTIMATED TYPE DISTRIBUTIONS

Type	Apparent from guesses	Econometric from guesses	Econometric from guesses, excluding random	Econometric from guesses, with specification test	Econometric from guesses and search, with specification test
L1	20	43	37	27	29
L2	12	20	20	17	14
L3	3	3	3	1	1
D1	0	5	3	1	0
D2	0	0	0	0	0
Eq.	8	14	13	11	10
Soph.	0	3	2	1	1
Unclassified	45	0	10	30	33

Note: The far-right-hand column includes 17 OB subjects classified by their econometric-from-guesses type estimates.

For those 45 subjects, there is some room for doubt about whether CGC's specification omits relevant types and/or overfits by including irrelevant types.

To test for this, CGC conducted a specification test, which suggests that the types estimated to be in the population are relevant and that any omitted types are at most 1-2% of the population, hence not worth modeling.

Lessons from the experiments for modeling strategic behavior

First, Nagel's 1995 *AER* subjects' initial guesses resembled neither equilibrium plus noise nor QRE for any reasonable distribution.

Nagel's results also suggest that even rationalizability is too strong: most subjects' guesses respected *k*–rationalizability only for small values of *k*.

Finally, Nagel's results call into question the common simplifying assumption that strategic thinking is homogeneous in the population.

No model that imposes homogeneity, as equilibrium plus noise and QRE do, will do full justice to subjects' behavior. Allowing heterogeneity of strategic thinking is essential for the explanations of Kahneman's Entry Magic, Yushchenko, and Lake Wobegon proposed below.

CGC's analysis significantly sharpens Nagel's conclusions, confirming by direct and econometric evidence and a specification test that a level-*k* model with a uniform random *L0* and only *L1*, *L2*, *L3*, and, possibly, *Equilibrium* subjects explains a large fraction of subjects' deviations from equilibrium in their games.

In particular:

- There are no *Dk* subjects. CGC's subjects respect iterated dominance to the extent that *Lk* types do, not because they explicitly perform it.

- Although level-*k* subjects make decisions that, via the iterated best responses that govern their strategic thinking) respect *k*-rationalizability, their presence is limited to small values of *k*, so even the *Lk* types respect *k*-rationalizability for at most small values of *k*.

- There are no *Sophisticated* subjects. Even the most sophisticated subjects seem to favor rules of thumb over less structured strategic thinking.

  (The jury is still out on the extent to which this conclusion generalizes.)

- CGC's evidence and analysis are more precise than previous studies of initial responses to normal-form games, but their conclusions are fully consistent with the results of earlier studies as well as folk game theory.

Illustration of level-*k* analyses of matrix games with unique mixed-strategy equilibria: M. M. Kaye's *The Far Pavilions*

I now give a simple example that illustrates applications of level-*k* models.

In M. M. Kaye's novel *The Far Pavilions*, the main male character, Ash, is trying to escape from his Pursuers along a North-South road.

Ash and his Pursuers have *strategically simultaneous* choices between North and South—although their choices are time-sequenced, the Pursuers must make their choice irrevocably before they learn Ash's choice.

If the Pursuers catch Ash, they gain 2 and he loses 2. But South is warm, and North is the Himalayas with winter coming. Thus both Ash and the Pursuers gain an extra 1 for choosing South, whether or not Ash is caught:

		Pursuers	
		South (*q*)	**North**
Ash	**South (*p*)**	**3** **-1**	**0** **1**
	North	**1** **0**	**2** **-2**

***Far Pavilions* Escape!**

Escape! has a unique equilibrium in mixed strategies, in which:

$$3p + 1(1 - p) = 0p + 2(1 - p) \text{ or } p = 1/4, \text{ and}$$

$$-1q + 1(1 - q) = 0q - 2(1 - q) \text{ or } q = 3/4.$$

This equilibrium responds to the payoff asymmetry between South and North in a decision-theoretically intuitive way for Pursuers (because $q = 3/4 >$ the 1/2 of equilibrium without the payoff asymmetry) but counterintuitively for Ash (because $p = 1/4 < 1/2$).

Although the equilibrium does not fully reflect intuition, experimental data from such games suggest that people's decisions often do reflect intuition.

E.g., Camerer reports informally gathered data for a perturbed Matching Pennies game (see also Rosenthal, Shachat, and Walker (2003 *IJGT*)):

	L (33%)	**R (67%)**
T (72%)	**2**, **0**	**0**, **1**
B (28%)	**0**, **1**	**1**, **0**

Perturbed Matching Pennies

The equilibrium mixed-strategy probabilities are Pr{T} = Pr{B} = 0.5 for Row and Pr{L} = 0.33 and Pr{R} = 0.67 for Column.

Although Column players are "right on" the equilibrium mixture, Row players overplay their superficially more attractive strategy T, not realizing that this allows a sophisticated Column to neutralize Row's advantage.

(Perhaps unsurprisingly, because that realization may require fixed-point reasoning.)

Meanwhile, back in the novel, Ash overcomes his fear of freezing and goes North. The Pursuers—unimaginatively—go South, Ash escapes, and the novel continues…

Meanwhile, back in the novel, Ash overcomes his fear of freezing and goes North. The Pursuers—unimaginatively—go South, Ash escapes, and the novel continues…romantically…

Meanwhile, back in the novel, Ash overcomes his fear of freezing and goes North. The Pursuers—unimaginatively—go South, Ash escapes, and the novel continues…romantically…for 900 more pages.

In equilibrium the observed outcome {Ash North, Pursuers South} has probability $(1 - p)q = 9/16$: a fit much better than random.

But try a level-*k* model with a uniformly random *L0*:

Types	**Ash**	**Pursuers**
L0	**uniformly random**	**uniformly random**
L1	**South**	**South**
L2	**North**	**South**
L3	**North**	**North**
L4	**South**	**North**
L5	**South**	**South**

Lk **types' decisions in *Far Pavilions* Escape!**

The level-*k* model precisely and correctly predicts the outcome provided that Ash is either *L2* or *L3* and the Pursuers are either *L1* or *L2*.

How do we know Ash's type (*L2* or *L3*)? One advantage of using fiction as data is that the narrative sometimes reveals cognition as well as decisions:

Ash's mentor—Koda Dad, played by Omar Sharif in the HBO miniseries—gives Ash the following advice (p. 97 of the novel):

"…ride hard for the north, since they will be sure you will go southward where the climate is kinder…").

Koda Dad's advice reflects the belief that the Pursuers think Ash is *L1*, so that Ash will go south because it's "kinder" and that (assuming that the Pursuers are uniform random *L0*) the Pursuers are no more likely to catch him there.

Thus Koda Dad must think the Pursuers are *L2*.

Hence Koda Dad advises Ash to think like an *L3*, and go North.

L3 ties my personal best *k* for a clearly explained level-*k* type in fiction. I suspect even postmodern fiction may have no *Lk*s higher than *L3*: they wouldn't be credible. I also doubt that one can find fixed-point reasoning.

Of course, most applications don’t come with an omniscient author identifying characters’ strategic thinking types for us.

But if the game is clearly defined and we have enough data, we can specify a level-*k* model, derive its implications, and use them to estimate the population frequency distribution of types and their precisions, as illustrated below.

Alternatively, we can calibrate the model using previous estimates for similar applications.

Returning to Camerer’s experiment, for example, an *L1* Row plays T and an *L1* Column plays L and R with equal probabilities (for logit or alternative payoff-driven error structures). An *L2* Row plays T and an *L2* Column plays R. An *L3* Row plays B and an *L3* Column plays R.

	L (33%)	**R (67%)**
T (72%)	**2**, **0**	**0**, **1**
B (28%)	**0**, **1**	**1**, **0**

Perturbed Matching Pennies

With a plausible mixture of 50% *L1*s, 30% *L2*s, and 20% *L3*s in both player roles—it’s natural to impose symmetry when roles are filled randomly from the same population—the level-*k* model’s predicted choice frequencies are 80% T for Row and 25% L for Column: Not a perfect fit, but reasonable.

The outcome resembles a “purified” mixed-strategy equilibrium.

But the level-*k* model predicts choice frequencies that deviate from the equilibrium probabilities for Row, Pr{T} = Pr{B} = 0.5, in the intuitive direction.

Similarly, in *Far Pavilions* Escape!, even though *Lk* types don't normally randomize, the heterogeneity of thinking reflected by the estimated distribution implies a mixture of decisions that reflects strategic uncertainty.

		Pursuers	
		South (q)	**North**
Ash	**South (p)**	**3** **-1**	**0** **1**
	North	**1** **0**	**2** **-2**

Far Pavilions **Escape!**

Suppose, for example, that each player role is filled from a 50-50 mixture of *L1*s and *L2*s and there are no errors.

Then Ash goes South with probability 0.5 > 1/4 (the equilibrium probability) and the Pursuers go South with probability 1 > 3/4 (the equilibrium probability).

Although the implied mixture of decisions again somewhat resembles a "purified" equilibrium, the model again deviates from equilibrium in the direction that intuition suggests: this time for both player roles.

Kahneman's Entry Magic: asymmetric coordination via structure in entry games

I now use a simple level-*k* model to suggest an explanation of Kahneman's Entry Magic.

In market-entry experiments, *n* subjects choose simultaneously between entering ("In") and staying out ("Out") of a market with given capacity.

For simplicity, assume that Out yields zero profit, no matter how many subjects enter.

In yields a given positive profit if no more subjects enter than a given market capacity; but a given negative profit if too many enter.

I will simplify Camerer, Ho, and Chong's (2004 *QJE*, Section III.C) CH analysis of *n*–person entry games to a level-*k* analysis of two-person Battle of the Sexes games, which are like two-person market-entry games with capacity one, and which makes the central points as simply as possible.

(Goldfarb and Yang (2008 *Journal of Marketing Research*) give a CH analysis of field data on analogous technology adoption games.)

Because players have no way to distinguish their symmetric roles, it is not sensible to predict systematic differences in behaviour across roles.

The natural equilibrium benchmark prediction is the symmetric mixed-strategy equilibrium, in which each player enters with a probability that makes all players indifferent between In and Out.

	In	**Out**
In	**0**, **0**	***a***, **1**
Out	**1**, ***a***	**0**, **0**

Battle of the Sexes

In Battle of the Sexes with $a > 1$, the unique symmetric equilibrium has

$$p \equiv \Pr\{\text{In}\} = a/(1+a) \text{ for both players.}$$

This mixed-strategy equilibrium yields an expected number of entrants roughly equal to market capacity ($2a/(1+a) \approx 1$, at least for a close to 1), but there is a positive probability that either too many or too few will enter.

With $p \equiv \Pr\{\text{In}\} = a/(1+a)$ for both players, the equilibrium expected coordination rate is $2p(1 - p) = 2a/(1+a)^2$.

Players' equilibrium expected payoffs are $a/(1+a)$, which is < 1 when $a > 1$: worse for each player than his worst pure-strategy equilibrium.

Even so, Kahneman's subjects regularly had better ex post coordination (number of entrants stochastically closer to market capacity) than in the symmetric equilibrium.

This led Kahneman to remark, "…to a psychologist, it looks like magic."

(But no one would be at all surprised by this unless he believed in equilibrium, so Kahneman should have said, "…to a *game theorist*, it looks like magic.")

Now consider a level-*k* model in which each player follows one of four types, *L1*, *L2*, *L3*, or *L4*, with each role filled by a draw from the same distribution.

Assume for simplicity that the frequency of *L0* is 0, and that *L0* chooses its action uniformly randomly, with Pr{In} = Pr{Out} = 1/2.

*L1*s mentally simulate *L0*s' random decisions and best respond, thus, with $a > 1$, choosing In; *L2*s choose Out; *L3*s choose In; and *L4*s choose Out.

	In	**Out**
In	**0** **0**	**1** ***a***
Out	***a*** **1**	**0** **0**

Battle of the Sexes

Type pairings	***L1***	***L2***	***L3***	***L4***
L1	**In, In**	**In, Out**	**In, In**	**In, Out**
L2	**Out, In**	**Out, Out**	**Out, In**	**Out, Out**
L3	**In, In**	**In, Out**	**In, In**	**In, Out**
L4	**Out, In**	**Out, Out**	**Out,In**	**Out, Out**

The predicted outcome distribution is determined by the outcomes of the possible type pairings and the type frequencies.

If both roles are filled from the same distribution, players have equal ex ante payoffs, proportional to the expected coordination rate.

L3 behaves like *L1*, and *L4* like *L2*. Lumping *L1* and *L3* together and letting v denote their total probability, and lumping *L2* and *L4* together, the expected coordination rate is $2v(1 - v)$.

This is maximized at $v = ½$, where it takes the value ½.

Thus for v near ½, which is behaviorally plausible, the coordination rate is close to ½. (For more extreme values the rate is worse, → 0 as $v \to 0$ or 1.)

By contrast, the mixed-strategy equilibrium expected coordination rate,
$2a/(1 + a)^2$, is maximized when $a = 1$, where it takes the value ½.

As $a \to \infty$, $2a/(1 + a)^2 \to 0$ like $1/a$. Even for moderate values of a, the level-k coordination rate is higher than the equilibrium rate.

The analysis illustrates the importance of the structured heterogeneity of strategic thinking a level-*k* model allows.

The level-*k* model, and the closely related CH model, yield a completely different view of asymmetric coordination via structure than a traditional refined-equilibrium model:

- Neither equilibrium nor refinements play any role in players' thinking.

- Coordination, when it occurs, is an accidental (though statistically predictable) by-product of players' non-equilibrium decision rules.

- Even though decisions are simultaneous and there is no communication or observation of the other's decision, the predictable heterogeneity of strategic thinking allows more sophisticated players such as *L2*s to mentally simulate the decisions of less sophisticated players such as *L1*s and accommodate them, just as Stackelberg followers would.

- This mental simulation doesn't work perfectly, because an *L2* is as likely to be paired with another *L2* as an *L1*. Neither would it work if strategic thinking were homogeneous. But it's very surprising that it works at all.

Yuschenko and Lake Wobegon: Framing effects in zero-sum two-person games

Consider Rubinstein, Tversky, and Heller’s 1993, 1996, 1998-99 (“RTH”) experiments with zero-sum, two-person “hide-and-seek” games with non-neutral framing of locations, analyzed by Crawford and Iriberri 2007 *AER*.

A typical seeker’s instructions (a hider’s instructions are analogous):

Your opponent has hidden a prize in one of four boxes arranged in a row. The boxes are marked as shown below: A, B, A, A. Your goal is, of course, to find the prize. His goal is that you will not find it. You are allowed to open only one box. Which box are you going to open?

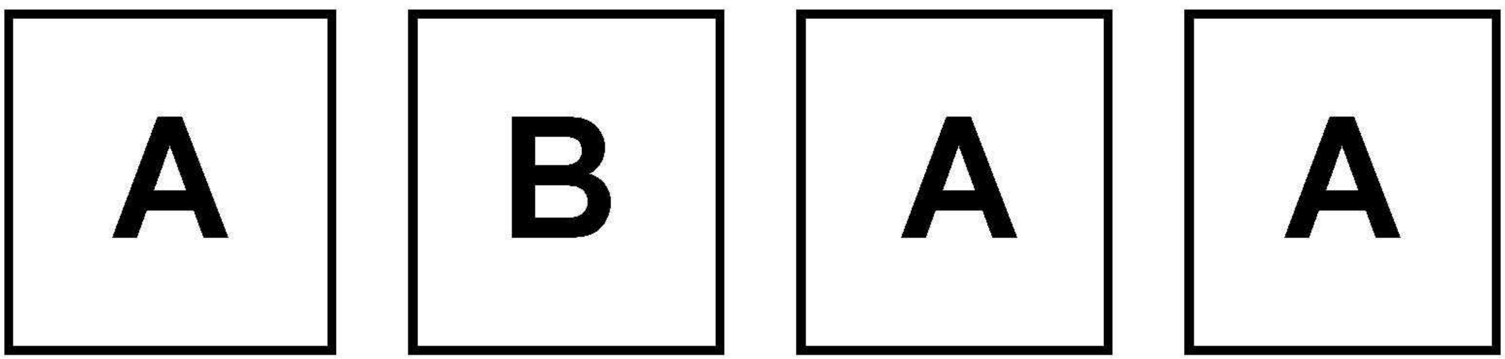

RTH’s framing of the hide-and-seek game is non-neutral in two ways:

● The “*B*” location is distinguished by its label.

● The two “*end A*” locations may be inherently focal.

This gives the “*central A*” location its own brand of uniqueness as the “least salient” location.

Mathematically this “negative” uniqueness is analogous to the “positive” uniqueness of “*B*”.

However, Crawford and Iriberri’s (2007 *AER*) analysis shows that its psychological effects are completely different.

RTH’s design is important as a tractable abstract model of a non-neutral cultural or geographic frame, or “landscape.”

Hide-and-seek games are often played on such landscapes, even though traditional game theory rules out any influence of the landscape by fiat.

This is well illustrated by the Yuschenko and Lake Wobegon quotations:

“Any government wanting to kill an opponent…would not try it at a meeting with government officials.”

“…in Lake Wobegon, the correct answer is usually ‘c’.”

Yuschenko’s meeting with government officials is analogous to RTH’s B, although in that example there’s nothing like RTH’s end locations.

With four possible choices arrayed left to right in the zero-sum game between a test designer deciding where to hide the correct answer and a clueless test-taker trying to guess where it is, the Lake Wobegon example is very close to RTH’s design.

RTH's hide-and-seek game has a clear equilibrium prediction, which leaves no room for framing to systematically influence the outcome.

The traditional theory of zero-sum two-person games is the strongpoint of noncooperative game theory, where the arguments for playing equilibrium strategies are immune to most of the usual counterarguments.

Yet framing has a strong and systematic effect in RTH's experiments, qualitatively the same around the world, with *Central A* (or its analogs in other treatments, as explained in the paper) most prevalent for hiders (37% in the aggregate) and even more prevalent for seekers (46%).

In this game any strategy, pure or mixed, is a best response to equilibrium beliefs. Thus one might argue that deviations do not violate the theory.

However, systematic deviations of aggregate choice frequencies from equilibrium probabilities must (with very high probability) have a cause that is partly common across players. They are therefore symptomatic of systematic deviations from equilibrium.

TABLE 1—AGGREGATE CHOICE FREQUENCIES IN RTH'S TREATMENTS

RTH-4	A	*B*	A	A
Hider (53; p = 0.0026)	9 percent	36 percent	40 percent	15 percent
Seeker (62; p = 0.0003)	13 percent	31 percent	45 percent	11 percent
RT-AABA–Treasure	A	A	*B*	A
Hider (189; p = 0.0096)	22 percent	35 percent	19 percent	25 percent
Seeker (85; p = 9E-07)	13 percent	51 percent	21 percent	15 percent
RT-AABA–Mine	A	A	*B*	A
Hider (132; p = 0.0012)	24 percent	39 percent	18 percent	18 percent
Seeker (73; p = 0.0523)	29 percent	36 percent	14 percent	22 percent
RT-1234–Treasure	1	*2*	3	4
Hider (187; p = 0.0036)	25 percent	22 percent	36 percent	18 percent
Seeker (84; p = 3E–05)	20 percent	18 percent	48 percent	14 percent
RT-1234–Mine	1	*2*	3	4
Hider (133; p = 6E-06)	18 percent	20 percent	44 percent	17 percent
Seeker (72; p = 0.149)	19 percent	25 percent	36 percent	19 percent
R-ABAA	A	*B*	A	A
Hider (50; p = 0.0186)	16 percent	18 percent	44 percent	22 percent
Seeker (64; p = 9E-07)	16 percent	19 percent	54 percent	11 percent

Notes: Sample sizes and p-values for significant differences from equilibrium in parentheses; salient labels in italics; order of presentation of locations to subjects as shown.

Crawford and Iriberri's Table 1

RTH's results raise several puzzles:

- Hiders' and seekers' responses are unlikely to be completely non-strategic in such simple games. So if they aren't following equilibrium logic, what are they doing?

- On average hiders are as smart as seekers, so hiders tempted to hide in *central A* should realize that seekers will be just as tempted to look there. Why do hiders allow seekers to find them 32% of the time when they could hold it down to 25% via the equilibrium mixed strategy?

- Further, why do seekers choose *central A* (or its analogs) even more often (46% in Table 3 below) than hiders (37%)?

Note that although the payoff structure of RTH's game is asymmetric, QRE ignores labeling and (logit or not) coincides with equilibrium in the game, and so does not help to explain the asymmetry of choice distributions.

The role asymmetry in subjects' behavior and how it is linked to the game's payoff asymmetry points strongly in the direction of a level-*k* or CH model, and is a mystery from the viewpoint of other theories I am aware of.

In constructing such a model, defining *L0* as uniform random would be unnatural, given the non-neutral framing of decisions and that *L0* describes others' instinctive responses.

(It would also make *Lk* the same as *Equilibrium* for $k > 0$.)

But a level-*k* model with a role-independent *L0* that probabilistically favors salient locations yields a simple explanation of RTH's results.

Assume that *L0* hiders and seekers both choose A, B, A, A with probabilities $p/2$, q, $1 - p - q$, $p/2$ respectively, with $p > ½$ and $q > ¼$.

L0 favors both the end locations and the B location, equally for hiders and seekers, but the model lets the data decide which is more salient.

For behaviorally plausible type distributions (estimated 0% *L0*, 19% *L1*, 32% *L2*, 24% *L3*, 25% *L4*—almost hump-shaped), a level-*k* model gracefully explains the major patterns in RTH's data, the prevalence of *central A* for hiders and its even greater prevalence for seekers:

● Given *L0*'s attraction to salient locations, *L1* hiders choose *central A* to avoid *L0* seekers and *L1* seekers avoid *central A* searching for *L0* hiders (the data suggest that end locations are more salient than B).

● For similar reasons, *L2* hiders choose *central A* with probability between 0 and 1 (breaking payoff ties randomly) and *L2* seekers choose it with probability 1.

● *L3* hiders avoid *central A* and *L3* seekers choose it with probability between zero and one (breaking payoff ties randomly).

● *L4* hiders and seekers both avoid *central A.*

TABLE 2—TYPES' EXPECTED PAYOFFS AND CHOICE PROBABILITIES IN RTH'S GAMES WHEN $p > 1/2$ AND $q > 1/4$

Hider	Expected payoff $p < 2q$	Choice probability $p < 2q$	Expected payoff $p > 2q$	Choice probability $p > 2q$	Seeker	Expected payoff $p < 2q$	Choice probability $p < 2q$	Expected payoff $p > 2q$	Choice probability $p > 2q$
L0 (Pr. r)					*L0* (Pr. r)				
A	–	$p/2$	–	$p/2$	A	–	$p/2$	–	$p/2$
B	–	q	–	q	B	–	q	–	q
A	–	$1-p-q$	–	$1-p-q$	A	–	$1-p-q$	–	$1-p-q$
A	–	$p/2$	–	$p/2$	A	–	$p/2$	–	$p/2$
L1 (Pr. s)					*L1* (Pr. s)				
A	$1-p/2 < 3/4$	0	$1-p/2 < 3/4$	0	A	$p/2 > 1/4$	0	$p/2 > 1/4$	1/2
B	$1-q < 3/4$	0	$1-q < 3/4$	0	B	$q > 1/4$	1	$q > 1/4$	0
A	$p+q > 3/4$	1	$p+q > 3/4$	1	A	$1-p-q < 1/4$	0	$1-p-q < 1/4$	0
A	$1-p/2 < 3/4$	0	$1-p/2 < 3/4$	0	A	$p/2 > 1/4$	0	$p/2 > 1/4$	1/2
L2 (Pr. t)					*L2* (Pr. t)				
A	1	1/3	1/2	0	A	0	0	0	0
B	0	0	1	1/2	B	0	0	0	0
A	1	1/3	1	1/2	A	1	1	1	1
A	1	1/3	1/2	0	A	0	0	0	0
L3 (Pr. u)					*L3* (Pr. u)				
A	1	1/3	1	1/3	A	1/3	1/3	0	0
B	1	1/3	1	1/3	B	0	0	1/2	1/2
A	0	0	0	0	A	1/3	1/3	1/2	1/2
A	1	1/3	1	1/3	A	1/3	1/3	0	0
L4 (Pr. v)					*L4* (Pr. v)				
A	2/3	0	1	1/2	A	1/3	1/3	1/3	1/3
B	1	1	1/2	0	B	1/3	1/3	1/3	1/3
A	2/3	0	1/2	0	A	0	0	0	0
A	2/3	0	1	1/2	A	1/3	1/3	1/3	1/3
Total	$p < 2q$		$p > 2q$		Total	$p < 2q$		$p > 2q$	
A	$rp/2+(1-\varepsilon)[t/3+u/3]+(1-r)\varepsilon/4$		$rp/2+(1-\varepsilon)[u/3+v/2]+(1-r)\varepsilon/4$		A	$rp/2+(1-\varepsilon)[u/3+v/3]+(1-r)\varepsilon/4$		$rp/2+(1-\varepsilon)[s/2+v/3]+(1-r)\varepsilon/4$	
B	$rq+(1-\varepsilon)[u/3+v]+(1-r)\varepsilon/4$		$rq+(1-\varepsilon)[t/2+u/3]+(1-r)\varepsilon/4$		B	$rq+(1-\varepsilon)[s+v/3]+(1-r)\varepsilon/4$		$rq+(1-\varepsilon)[u/2+v/3]+(1-r)\varepsilon/4$	
A	$r(1-p-q)+(1-\varepsilon)[s+t/3]+(1-r)\varepsilon/4$		$r(1-p-q)+(1-\varepsilon)[s+t/2]+(1-r)\varepsilon/4$		A	$r(1-p-q)+(1-\varepsilon)[t+u/3]+(1-r)\varepsilon/4$		$r(1-p-q)+(1-\varepsilon)[t+u/2]+(1-r)\varepsilon/4$	
A	$rp/2+(1-\varepsilon)[t/3+u/3]+(1-r)\varepsilon/4$		$rp/2+(1-\varepsilon)[u/3+v/2]+(1-r)\varepsilon/4$		A	$rp/2+(1-\varepsilon)[u/3+v/3]+(1-r)\varepsilon/4$		$rp/2+(1-\varepsilon)[s/2+v/3]+(1-r)\varepsilon/4$	

TABLE 3—PARAMETER ESTIMATES AND LIKELIHOODS FOR THE LEADING MODELS IN RTH'S GAMES

Model	Ln L	Parameter estimates	Player	A	B	A	A	MSE
			Observed or predicted choice frequencies					
Observed frequencies (624 hiders, 560 seekers)			H S	0.2163 0.1821	0.2115 0.2054	0.3654 0.4589	0.2067 0.1536	–
Equilibrium without perturbations	−1641.4		H S	0.2500 0.2500	0.2500 0.2500	0.2500 0.2500	0.2500 0.2500	0.00970
Equilibrium with restricted perturbations	−1568.5	$e_H \equiv e_S = 0.2187$ $f_H \equiv f_S = 0.2010$	H S	0.1897 0.1897	0.2085 0.2085	0.4122 0.4122	0.1897 0.1897	0.00084
Equilibrium with unrestricted perturbations	−1562.4	$e_H = 0.2910, f_H = 0.2535$ $e_S = 0.1539, f_S = 0.1539$	H S	0.2115 0.1679	0.2115 0.2054	0.3654 0.4590	0.2115 0.1679	0.00006
Level-*k* with a role-symmetric *L0* that favors salience	−1564.4	$p > 1/2$ and $q > 1/4, p > 2q$, $r = 0, s = 0.1896, t = 0.3185$, $u = 0.2446, v = 0.2473, \varepsilon = 0$	H S	0.2052 0.1772	0.2408 0.2047	0.3488 0.4408	0.2052 0.1772	0.00027
Level-*k* with a role-asymmetric *L0* that favors salience for seekers and avoids it for hiders	−1563.8	$p_H < 1/2$ and $q_H < 1/4$, $p_S > 1/2$ and $q_S > 1/4$, $r = 0, s = 0.66, t = 0.34$, $\varepsilon = 0.72$; $u \equiv v \equiv 0$ imposed	H S	0.2117 0.1800	0.2117 0.1800	0.3648 0.4600	0.2117 0.1800	0.00017
Level-*k* with a role-symmetric *L0* that avoids salience	−1562.5	$p < 1/2$ and $q < 1/4, p < 2q$, $r = 0, s = 0.3636, t = 0.0944$, $u = 0.3594, v = 0.1826, \varepsilon = 0$	H S	0.2133 0.1670	0.2112 0.2111	0.3623 0.4549	0.2133 0.1670	0.00006

Crawford and Iriberri's Table 3

Note that only a heterogeneous population with substantial frequencies of *L2* and *L3* as well as *L1* (estimated 0% *L0*, 19% *L1*, 32% *L2*, 24% *L3*, 25% *L4*) can reproduce the aggregate patterns in the data.

(Even though there is a nonnegligible estimated frequency of *L4*s, they don't really matter here because they never choose *central A* (Table 2 above), hence they are not implicated in the major aggregate patterns.

For the same reason, their frequency is not well identified in the estimation.)

For example, Crawford and Iriberri estimate (Table 3 above, row 5) that the salience of an end location is greater than the salience of the *B* ($p > 2q$).

Given this, a 50-50 mix of *L1*s and *L2*s in both player roles would imply (Table 2 above, right-most columns in each panel) 75% of hiders but only 50% of seekers choosing *central A*, in contrast to the 37% of hiders and 46% of seekers who did choose *central A.*

In Crawford and Iriberri's analysis of RTH's data, the role asymmetry in aggregate behavior follows naturally from the asymmetry of the game's payoff structure, via hiders' and seekers' asymmetric responses to *L0*'s *role-symmetric* choices.

Allowing *L0* to vary across roles as in Bacharach and Stahl 2000 *GEB*, although it yields a small improvement in fit (Table 3), would beg the question of why subjects' responses were so role-asymmetric.

Crawford and Iriberri's analysis also suggests that allowing *L0* to vary across roles leads to overfitting.

RTH took the main patterns in their data as evidence that their subjects did not think strategically:

● "The finding that both choosers and guessers selected the least salient alternative suggests little or no strategic thinking."

● "In the competitive games, however, the players employed a naïve strategy (avoiding the endpoints), that is not guided by valid strategic reasoning. In particular, the hiders in this experiment either did not expect that the seekers too, will tend to avoid the endpoints, or else did not appreciate the strategic consequences of this expectation."

RTH could have said the same thing about the Yuschenko quotation:

- “Any government wanting to kill an opponent…would not try it at a meeting with government officials”,

to which a game theorist would (almost involuntarily) respond:

- “If that’s what people think, a meeting with government officials is exactly where *I* would try to poison Yushchenko.”

But strategic thinking need not be equilibrium thinking.

Crawford and Iriberri’s analysis suggests that RTH’s subjects were actually quite strategic and in fact more than usually sophisticated (with many *L3*s and even some *L4*s, even though in most settings *L1*s and *L2*s are more common)—they just didn’t follow equilibrium logic.

Crawford and Iriberri’s analysis suggests that the Yushchenko quotation simply reflects the reasoning of an *L1* poisoner, or equivalently of an *L2* investigator reasoning about an *L1* poisoner.

Crawford and Iriberri tested for portability by using the leading alternative models, estimated from RTH's data, to "predict" subjects' initial responses in the closest relative of RTH's games in the literature, O'Neill's (1987 *PNAS*) famous card-matching game.

O'Neill's game raises the same kinds of strategic issues as RTH's games, but with more complex patterns of wins and losses and different framing.

In O'Neill's card-matching game, players simultaneously and independently choose one of four cards: A, 2, 3, J.

One player, say the row player—but the game was presented to subjects as a story, not a matrix—wins if there is a match on J or a mismatch on A, 2, or 3; the other player wins in the other cases.

	A	**2**	**3**	**J**
A	**1** **0**	**0** **1**	**0** **1**	**1** **0**
2	**0** **1**	**1** **0**	**0** **1**	**1** **0**
3	**0** **1**	**0** **1**	**1** **0**	**1** **0**
J	**1** **0**	**1** **0**	**1** **0**	**0** **1**

O'Neill's card-matching game

O’Neill’s game is like a hide-and-seek game, except that each player is a hider (h) for some locations and a seeker (s) for others.

A, 2, and 3 are strategically symmetric, and equilibrium (without payoff perturbations) has Pr{A} = Pr{2} = Pr{3} = 0.2, Pr{J} = 0.4.

	A (s)	**2 (s)**	**3 (s)**	**J (h)**
A (h)	**1** **0**	**0** **1**	**0** **1**	**1** **0**
2 (h)	**0** **1**	**1** **0**	**0** **1**	**1** **0**
3 (h)	**0** **1**	**0** **1**	**1** **0**	**1** **0**
J (s)	**1** **0**	**1** **0**	**1** **0**	**0** **1**

O’Neill’s card-matching game

The portability test directly addresses the issue of whether level-*k* models allow the modeler too much flexibility.

With regard to the flexibility of *L0*, first consider how to adapt our "psychological" specification of *L0* from RTH's to O'Neill's game.

"Anyone" should agree on the right kind of *L0*:

- A and J, "face" cards and end locations, are more salient than 2 and 3, but the specification should allow either A or J to be more salient.

That the RTH estimates suggested that their end locations are more salient than the *B* label does *not* dictate whether A or J is more salient, though it does reinforce that they are both more salient than 2 and 3.

This is a psychological issue, but because it is "only" a psychological issue, it is easy to gather evidence on it from different settings, and such evidence is more likely to yield convergence than if it were partly a strategic issue.

Further, because all that matters about *L0* is what it makes *L1*s do in each role, the remaining freedom to choose *L0* allows only two models.

With regard to the flexibility of the type frequencies, empirically plausible frequencies often imply severe limits on what decision patterns a level-*k* model can generate.

Readers of the first version of Crawford and Iriberri (2007 *AER*) often asked if the model could explain behavior in games other than RTH's.

O'Neill's game was the most natural choice in the experimental literature.

We did not have his data, but discussions of it (e.g. McKelvey and Palfrey (1995 *GEB*)) had been dominated by an "Ace effect": aggregated over all 105 rounds, row and column players played A with frequencies 22.0% and 22.6%, significantly above the equilibrium 20%.

(O'Neill speculated that this was because "…players were attracted by the powerful connotations of an Ace".

But—we thought—what about the equally powerful connotations of the Joker and its unique payoff role? They seem to make Joker even more salient than Ace, but in the aggregate data row subjects chose Joker with frequencies of only 36%, and column subjects with frequencies of only 43%.)

We also knew that with a plausible specification of *L0* and the resulting types' decisions in O'Neill's game (Tables A3 and A4 from the paper's web appendix, reproduced on the next two slides), no behaviorally plausible level-*k* model could make a row player ("Player 1") play A more than the equilibrium 20%:

Tables A3 and A4 show that, excluding *L0*s (which normally have 0 estimated frequencies) and restricting attention to Player 1, when A is more salient ($3j - a < 1$) only *L4* chooses A, and that with probability at most 1/3 (Table A3); and that when A is less salient ($3j - a > 1$) only *L3* chooses A, and that with probability at most 1/3 (Table A4).

This is *logically* possible, but in the first case it would require a population of 60% or more *L4*s, and in the second case it would require 60% or more *L3*s: in each case behaviorally extremely unlikely on the available evidence.

Thus, despite the flexibility of the estimated type distribution, the level-*k* model's structure and the principles that guide the specification of *L0* imply a strong restriction: that row players play A less than the equilibrium 20%.

Table A3. Types' Expected Payoffs and Choice Probabilities in O'Neill's Game when $3j - a < 1$

Player 1	Exp. Payoff $A+2j<1$	Choice Pr. $a+2j<1$	Exp. Payoff $a+2j>1$	Choice Pr. $a+2j>1$	Player 2	Exp. Payoff $a+2j<1$	Choice Pr. $a+2j<1$	Exp. Payoff $a+2j>1$	Choice Pr. $a+2j>1$
***L0* (Pr. *R*)**					***L0* (Pr. *r*)**				
A	-	a	-	A	**A**	-	a	-	a
2	-	$(1-a-j)/2$	-	$(1-a-j)/2$	**2**	-	$(1-a-j)/2$	-	$(1-a-j)/2$
3	-	$(1-a-j)/2$	-	$(1-a-j)/2$	**3**	-	$(1-a-j)/2$	-	$(1-a-j)/2$
J	-	j	-	J	**J**	-	j	-	j
***L1* (Pr. *s*)**					***L1* (Pr. *s*)**				
A	$1-a-j$	0	$1-a-j$	0	**A**	$a+j$	0	$a+j$	1
2	$(1+a-j)/2$	1/2	$(1+a-j)/2$	1/2	**2**	$(1-a+j)/2$	0	$(1-a+j)/2$	0
3	$(1+a-j)/2$	1/2	$(1+a-j)/2$	1/2	**3**	$(1-a+j)/2$	0	$(1-a+j)/2$	0
J	J	0	J	0	**J**	$1-j$	1	$1-j$	0
***L2* (Pr. *t*)**					***L2* (Pr. *t*)**				
A	0	0	0	0	**A**	0	0	0	0
2	0	0	1	1/2	**2**	½	0	1/2	0
3	0	0	1	1/2	**3**	½	0	1/2	0
J	1	1	0	0	**J**	1	1	1	1
***L3* (Pr. *u*)**					***L3* (Pr. *u*)**				
A	0	0	0	0	**A**	1	1/3	0	0
2	0	0	0	0	**2**	1	1/3	1/2	0
3	0	0	0	0	**3**	1	1/3	1/2	0
J	1	1	1	1	**J**	0	0	1	1
***L4* (Pr. *v*)**					***L4* (Pr. *v*)**				
A	2/3	1/3	0	0	**A**	1	1/3	1	1/3
2	2/3	1/3	0	0	**2**	1	1/3	1	1/3
3	2/3	1/3	0	0	**3**	1	1/3	1	1/3
J	0	0	1	1	**J**	0	0	0	0
Total	$a+2j<1$		$a+2j>1$		**Total**	$a+2j<1$		$a+2j>1$	
A	$ra+(1-\varepsilon)[v/3]+(1-r)\,\varepsilon/4$		$ra+(1-r)\,\varepsilon/4$		**A**	$ra+(1-\varepsilon)\,[u/3+v/3]+(1-r)\,\varepsilon/4$		$ra+(1-\varepsilon)\,[s+v/3]+(1-r)\,\varepsilon/4$	
2	$r(1-a-j)/2+(1-\varepsilon)\,[s/2+v/3]+(1-r)\,\varepsilon/4$		$r(1-a-j)/2+(1-\varepsilon)\,[s/2+t/2]+(1-r)\,\varepsilon/4$		**2**	$r(1-a-j)/2+(1-\varepsilon)\,[u/3+v/3]+(1-r)\,\varepsilon/4$		$r(1-a-j)/2+(1-\varepsilon)\,[v/3]+(1-r)\,\varepsilon/4$	
3	$r(1-a-j)/2+(1-\varepsilon)\,[s/3+v/3]+(1-r)\,\varepsilon/4$		$r(1-a-j)/2+(1-\varepsilon)\,[s/2+t/2]+(1-r)\,\varepsilon/4$		**3**	$r(1-a-j)/2+(1-\varepsilon)\,[u/3+v/3]+(1-r)\,\varepsilon/4$		$r(1-a-j)/2+(1-\varepsilon)\,[v/3]+(1-r)\,\varepsilon/4$	
J	$Rj+(1-\varepsilon)\,[t+u]+(1-r)\,\varepsilon/4$		$rj+(1-\varepsilon)\,[u+v]+(1-r)\,\varepsilon/4$		**J**	$rj+(1-\varepsilon)\,[s+t]+(1-r)\,\varepsilon/4$		$rj+(1-\varepsilon)\,[t+u]+(1-r)\,\varepsilon/4$	

Table A4. Types' Expected Payoffs and Choice Probabilities in O'Neill's Game when $3j - a > 1$

Player 1	Exp. Payoff	Choice Pr.	Player 2	Exp. Payoff	Choice Pr.
L0 (Pr. *R*)			*L0* (Pr. *r*)		
A	-	a	A	-	a
2	-	$(1-a-j)/2$	2	-	$(1-a-j)/2$
3	-	$(1-a-j)/2$	3	-	$(1-a-j)/2$
J	-	j	J	-	j
L1 (Pr. *S*)			*L1* (Pr. *s*)		
A	$1-a-j$	0	A	$a+j$	1
2	$(1+a-j)/2$	0	2	$(1-a+j)/2$	0
3	$(1+a-j)/2$	0	3	$(1-a+j)/2$	0
J	j	1	J	$1-j$	0
L2 (Pr. *T*)			*L2* (Pr. *t*)		
A	0	0	A	1	1/3
2	1	1/2	2	1	1/3
3	1	1/2	3	1	1/3
J	0	0	J	0	0
L3 (Pr. *U*)			*L3* (Pr. *u*)		
A	2/3	1/3	A	0	0
2	2/3	1/3	2	1/2	0
3	2/3	1/3	3	1/2	0
J	0	0	J	1	1
L4 (Pr. *V*)			*L4* (Pr. *v*)		
A	0	0	A	1/3	0
2	0	0	2	1/3	0
3	0	0	3	1/3	0
J	1	1	J	1	1
Total			**Total**		
A	$Ra+(1-\varepsilon)[u/3]+(1-r)\,\varepsilon/4$		A	$ra+(1-\varepsilon)[s+t/3]+(1-r)\,\varepsilon/4$	
2	$r(1-a-j)/2+(1-\varepsilon)[t/2+u/3]+(1-r)\,\varepsilon/4$		2	$r(1-a-j)/2+(1-\varepsilon)[t/3]+(1-r)\,\varepsilon/4$	
3	$R(1-a-j)/2+(1-\varepsilon)[t/2+u/3]+(1-r)\,\varepsilon/4$		3	$r(1-a-j)/2+(1-\varepsilon)[t/3]+(1-r)\,\varepsilon/4$	
J	$Rj+(1-\varepsilon)[s+v]+(1-r)\,\varepsilon/4$		J	$rj+(1-\varepsilon)[u+v]+(1-r)\,\varepsilon/4$	

Despite our fear that our explanation of RTH's results would be discredited by O'Neill's Ace Effect, we decided to order and analyze O'Neill's data, speculating, based on the level-*k* model's success in RTH's and other games, that his subjects' *initial* responses must not have had an Ace effect.

The initial responses were:

● 8% A, 24% 2, 12% 3, 56% J for rows, and

● 16% A, 12% 2, 8% 3, 64% J for columns.

No Ace effect!

On the contrary, for initial responses there was a huge Joker effect, an order of magnitude stronger than the Ace effect in the time-aggregated data (but never before mentioned in the literature).

(An order of magnitude stronger because (56 - 40)% and (64 - 40)% are respectively roughly ten times larger than (22 - 20)% and (22.6 - 20)%.)

Unlike the putative Ace effect, the actual Joker effect (and the other frequencies) *can* be explained by a level-*k* model with a plausible *L0* that probabilistically favors the salient A and J cards.

(The analysis also suggests that the Ace effect in the time-aggregated data was an accidental by-product of how subjects learned, not of salience at all.)

TABLE 5—COMPARISON OF THE LEADING MODELS IN O'NEILL'S GAME

Model	Parameter estimates	Player	A	2	3	J	MSE
Observed frequencies (25 Player 1s, 25 Player 2s)		1	0.0800	0.2400	0.1200	0.5600	–
		2	0.1600	0.1200	0.0800	0.6400	–
Equilibrium without perturbations		1	0.2000	0.2000	0.2000	0.4000	0.0120
		2	0.2000	0.2000	0.2000	0.4000	0.0200
Level-*k* with a role-symmetric *L0* that favors salience	$a > 1/4$ and $j > 1/4$ $3j - a < 1, a + 2j < 1$	1	0.0824	0.1772	0.1772	0.5631	0.0018
		2	0.1640	0.1640	0.1640	0.5081	0.0066
Level-*k* with a role-symmetric *L0* that favors salience	$a > 1/4$ and $j > 1/4$ $3j - a < 1, a + 2j > 1$	1	0.0000	0.2541	0.2541	0.4919	0.0073
		2	0.2720	0.0824	0.0824	0.5631	0.0050
Level-*k* with a role-symmetric *L0* that avoids salience	$a < 1/4$ and $j < 1/4$	1	0.4245	0.1807	0.1807	0.2142	0.0614
		2	0.1670	0.1807	0.1807	0.4717	0.0105
Level-*k* with a role-asymmetric *L0* that favors salience for locations for which player is a seeker and avoids it for locations for which player is a hider	$a_1 < 1/4, j_1 > 1/4;$ $a_2 > 1/4, j_2 < 1/4$ $3j_1 - a_1 < 1,$ $a_1 + 2j_1 < 1, 3a_2 + j_2 > 1$	1	0.1804	0.2729	0.2729	0.2739	0.0291
		2	0.1804	0.1804	0.1804	0.4589	0.0117

Note: in the printed table, "Observed or predicted choice frequencies" spans the Player, A, 2, 3 and J columns.

Crawford and Iriberri's Table 5

Importantly, Crawford and Iriberri's analysis traces the portability of the level-*k* model (in contrast to the alternative explanations considered in the paper) to the fact that *L0* is psychological rather than strategic, and that it is based on simple and universal intuition and evidence.

If *L0* were strategic, it would interact with the strategic structure in new ways in each new game, and it would be a rare event when one could extrapolate a specification from one game to another as Crawford and Iriberri did from RTH's games to O'Neill's.

Thus, the definition of *L0* as an instinctive, nonstrategic response is more that a convenient cognitive categorization: it is important for portability.

Adaptive Learning

Experimental evidence strongly suggests that whenever people have opportunities to observe other people's decisions in analogous games, strategic thinking is eclipsed by adaptive learning.

Learning models describe how players adjust their decisions over time in response to experience.

The learning process is usually modeled as repetition of a fixed “stage game,” so that the analogies are perfect.

The stage 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*) game discussed above).

Players view decisions in the stage game as the objects of choice, and the dynamics of their decisions are modeled (either directly, or indirectly in terms of their beliefs) as adjusting in a direction that would increase payoffs, other things equal, given the current state of the system.

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.

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 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).

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

Your choice of X		Median value of X chosen						
		7	6	5	4	3	2	1
	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

Your Choice of X		Smallest Value of X Chosen						
		7	6	5	4	3	2	1
	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.

The games are also closely related to Larry Summers's Bank Runs example:
"A crude but simple game, related to Douglas Diamond and Philip Dybvig's (1983 *JPE*) celebrated analysis of bank runs, illustrates some of the issues involved here. Imagine that everyone who has invested \$10 with me can expect to earn \$1, assuming that I stay solvent. Suppose that if I go bankrupt, investors who remain lose their whole \$10 investment, but that an investor who withdraws today neither gains nor loses. What would you do? Each individual judgment would presumably depend on one's assessment of my prospects, but this in turn depends on the collective judgment of all of the investors.

Suppose, first, that my foreign reserves, ability to mobilize resources, and economic strength are so limited that if any investor withdraws I will go bankrupt. It would be a Nash equilibrium (indeed, a Pareto-dominant one) for everyone to remain, but (I expect) not an attainable one. Someone would reason that someone else would decide to be cautious and withdraw, or at least that someone would reason that someone would reason that someone would withdraw, and so forth. This…would likely lead to large-scale withdrawals, and I would go bankrupt. It would not be a close-run thing. …Keynes's beauty contest captures a similar idea.

Now suppose that my fundamental situation were such that everyone would be paid off as long as no more than one-third of the investors chose to withdraw. What would you do then? Again, there are multiple equilibria: everyone should stay if everyone else does, and everyone should pull out if everyone else does, but the more favorable equilibria seems much more robust."—Lawrence Summers, "International Financial Crises: Causes, Prevention, and Cures," (2000 *AER*).

An *n*-person coordination game with Pareto-ranked equilibria. Summers presumes that some equilibrium will emerge, but his model of the influence of fragility on equilibrium selection may implicitly invoke initial responses to shocks followed by adaptive learning (although he cites Morris and Shin's (1998 *AER*) non-adaptive "global games" analysis).

The game Summers describes can be represented by a payoff table as follows:

		Summary statistic	
		In	**Out**
Representative player	**In**	**1**	**-10**
	Out	**0**	**0**

Bank Runs

The summary statistic is a measure of whether or not the required number of investors stays In.

In Summers's first example, all investors must stay In to prevent the bank from collapsing, so the summary statistic takes the value In if and only if all (but the representative player) stay In.

In his second example two-thirds of the investors need to stay In, so the summary statistic takes the value In if and only if (adding in the representative player) this is the case.

In each example there are two pure-strategy equilibria: "all-In" and "all-Out".

(There is also a mixed-strategy equilibrium in which the probability that the summary statistic equals In just balances the benefits of In and Out; but this equilibrium is behaviorally implausible.)

Aside: Initial responses in Summers's Bank Runs game

What happens depends on players' initial responses to the game as shaped by their strategic thinking: which equilibrium's basin of attraction, “all-In” or “all-Out”, the initial responses fall into.

The leading models of initial responses for games like this in traditional game theory are Harsanyi and Selten's notions of payoff- and risk-dominance.

Payoff-dominance favors equilibria that are Pareto-superior to other equilibria, hence here uniquely favors the all-In equilibrium, for any value of the population size *n* and any value of the deviation costs which here equals -10. This seems behaviorally quite unlikely even for small *n* and small -10.

The basic idea of risk-dominance (the precise formalization is controversial) is to choose the equilibrium with the largest “basin of attraction” in beliefs space.

In 2x2 symmetric two-person games, this amounts to selecting the equilibrium that results if each player best responds to a uniform random prior over the other’s strategies (just as *L1* does when *L0* is uniform random).

Thus for population size 2, risk-dominance favors the all-Out equilibrium.

In 2x2 symmetric games for population *n* > 2, risk-dominance again favors the equilibrium with the larger basin of attraction in beliefs space. Assuming independence, with Summers’s payoffs risk-dominance favors the all-Out equilibrium for any *n* > 2, even if only two-thirds need to stay In.

Now consider a level-*k* model.

In a realistic application *L0* would have to reflect market psychology, but to illustrate how the model works, I assume a uniform random *L0*.

In *n*-person games it is also possible to define a level-*k* model in which *L0* is correlated across players instead of independent.

(Risk-dominance is usually defined assuming independence, but correlation is possible there too. Correlation is irrelevant in defining payoff-dominance.)

In Summers's first example, where the summary statistic takes the value In only when all stay In, *L1*'s decision is Out with either independent or correlated *L0*.

In Summers's second example, where the summary statistic takes the value In when two-thirds or more stay In, *L1*'s decision is still Out in either case.

In all cases *L2* and higher types also stay Out, so if the frequency of *L0* is 0, the outcome is observationally equivalent to the all-Out equilibrium.

Now consider an example like Bank Runs in which the summary statistic takes the value In when *one-third* or more of the investors stay In.

If, say, $n = 6$, then given a choice of In by the representative player himself, the summary statistic will be In unless all five other players stay Out.

If *L0* is independent, *L1* assigns all others staying Out probability $1/2^5 \approx 0.03$.

If *L0* is correlated, *L1* assigns all others staying Out probability ½.

In the former case, *L1* and therefore all higher *Lk* types stay In, and the outcome is observationally equivalent to the all-In equilibrium.

In the latter case, *L1* and therefore all higher *Lk* types stay Out, and the outcome is observationally equivalent to the all-Out equilibrium.

In each of these symmetric coordination games, the level-*k* model derives the outcome from strategic responses to instinctive reactions to the game.

Unlike traditional coordination refinements, the level-*k* approach is easy to combine with richer models of market psychology, via *L0*.

And because such an *L0* is a psychological rather than a strategic concept, it is easier to extrapolate its specification across games, as illustrated below.

Again neither equilibrium nor refinements play any role in players' thinking.

And coordination, when it occurs, is again an accidental by-product of players' non-equilibrium, level-*k* decision rules.

Because in these symmetric coordination games *L1* responses to a uniform random *L0* are in equilibrium, there is no “magic”:

The level-*k* model reduces to an equilibrium selection device, which coincides here with risk-dominance, but need not do so in general.

In 2×2 symmetric coordination games *L1* responses to a uniform random *L0* also coincide with the equilibrium selected by a global games analysis.

Selecting an equilibrium via *L1* responses seems empirically more promising, because *L1* responses are less cognitively taxing and are directly suggested by experimental evidence.

End of aside

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	5	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

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	0	0
# of 5s	6	8	2	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	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 $FI(P)$ and the complement of $FI(P)$.

Explaining VHBB’s 1990 and 1991 results

"Rational learning"

One possible source of explanations of VHBB's results is a rational learning model, which models the learning process as an equilibrium in the repeated game that describes the entire learning process.

Rational learning is unhelpful in explaining VHBB’s 1990 and 1991 results because any pattern of perfectly coordinated jumping from one pure-strategy equilibrium to another over time is a rational learning equilibrium.

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*).

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.

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.

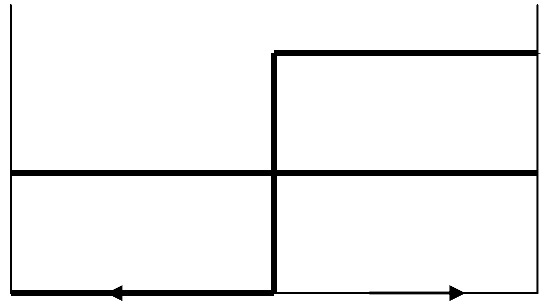

By symmetry, random shocks are neutral, equally likely to flip the population from all-Stag to all-Rabbit or vice versa.

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.

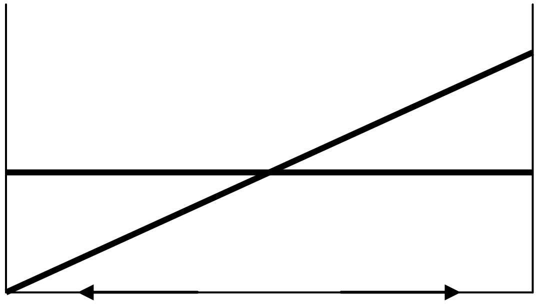

Random shocks are again neutral; and with random initial conditions, the population would be equally likely to converge to all-Stag or all-Rabbit.

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.

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.

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 less efficient in larger groups (holding the order statistic constant, measured from the bottom) because it requires coherence among more independent decisions.

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.